

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ROKU, INC.,

Petitioner,

v.

FLEXIWORLD TECHNOLOGIES, INC.,

Patent Owner.

PTAB Case No. PGR2021-00112

Patent No. 11,029,903

**DECLARATION OF SAMRAT BHATTACHARJEE, PH.D.
IN SUPPORT OF PETITION FOR POST GRANT REVIEW OF
U.S. PATENT NO. 11,029,903**

Table of Contents

	Page
I. INTRODUCTION AND SUMMARY OF TESTIMONY	1
II. OVERVIEW OF THE TECHNOLOGY	5
1. Wireless Communications	5
B. The '903 Patent	10
1. The '903 Patent Purports to Eliminate the Need to Pre- Install Drivers for Printers and Other Output Devices.	10
2. The '903 Patent's Pervasive Output Process	12
C. Claim Construction	19
III. UNPATENTABILITY OF THE CHALLENGED PATENT CLAIMS.....	19
A. Relevant Legal Standards.....	19
1. Lack of Written Description Under 35 U.S.C. § 112	19
2. Subject Matter Eligibility Under 35 U.S.C. § 101.....	20
B. Lack of Written Description.....	20
1. The Written Description Fails to Support the Independent Claims.	21
a. The '903 Patent's Written Description Does Not Support an Output System / Device that Connects to Servers Over the Internet.....	21
b. The '903 Patent's Written Description Does Not Support an Output System / Device that Sends the Job Object and Document / Content Object to the Server.	27
c. The '903 Patent's Written Description Does Not Support an Output System / Device that Receives Indications of Content Selections From the User.....	31
d. The '903 Patent's Written Description Does Not Support an Output System / Device that Receives Output Data From the Server.....	35
2. The Written Description Fails to Support the Dependent Claims for the Same Reasons.	38

a.	Dependent Claims 3, 14, and 18.....	38
3.	Information Apparatus 100 is not an “Output System” (claims 1 & 15) or “Wireless Output Device” (claim 8).	40
C.	Subject Matter Ineligibility	46
1.	Claim 1	47
a.	Preamble	47
b.	Claim 1 Recites Conventional Computing Devices	48
(i)	“Output System”	48
(ii)	“One or More Servers”.....	56
(iii)	“Client Device”	59
c.	Claim 1 Recites Conventional Wireless Communications	60
d.	Claim 1 Recites Conventional Steps	60
(i)	Step 1: Obtaining Authentication Information.....	61
(ii)	Step 2: Connecting to a Wireless LAN	65
(iii)	Step 3: Connecting to a Server on the Internet.....	67
(iv)	Step 4: Sending Authentication Information	70
(v)	Step 5: Accessing a Content Service.....	74
(vi)	Step 6: Receiving User Selection of Content.....	75
(vii)	Step 7: Identifying Selected Content to the Server.....	79
(viii)	Step 8: Receiving Content.....	81
(ix)	Step 9: Processing Content.....	85
(x)	Step 10: Delivering Content to an Output Device.....	92
(xi)	Step 11: Executing a Discovery Operation	93
(xii)	Step 12: Receiving Content from a Client Device.....	95
e.	The Claim Elements as an Ordered Combination	98

2.	Independent Claims 8 and 15.....	100
3.	Dependent Claims	106
	(i) Claims 2, 13 and 16.....	106
	(ii) Claims 3, 14, and 18.....	109
	(iii) Claims 4, 11, and 19.....	110
	(iv) Claim 5	111
	(v) Claims 6, 7 and 20.....	113
	(vi) Claim 9	114
	(vii) Claims 10 and 17.....	117
	(viii) Claim 12	118
IV.	CONCLUSION.....	121

LIST OF APPENDICES

Appendix A	Curriculum Vitae
Appendix B	Excerpts from Jim Geier, <i>Wireless LANs: Implementing Interoperable Networks</i> (MacMillan, 1999)
Appendix C	John Markoff, <i>New Economy: Airborne and grass roots. By popular acclaim, a wireless format with a name only a geek could love is taking hold.</i> (New York Times, Sec. C, p. 5, October 30, 2000)
Appendix D	Jaap C. Haartsen, The Bluetooth Radio System (IEEE Personal Communications, February 2000)
Appendix E	Golden G. Richard III, Service Advertisement and Discovery: Enabling Universal Device Cooperation (IEEE Internet Computing, September / October 2000)
Appendix F	Excerpt from Tom Sheldon, <i>McGraw Hill Encyclopedia of Networking & Telecommunications</i> (Osborne / McGraw Hill, 2001) at pp. 1131-33 (“Service Advertising and Discovery”).
Appendix G	Charlie Russel and Sharon Crawford, <i>Running Microsoft Windows NT Server 4.0</i> (1997)
Appendix H	Excerpts from Lon Poole & John Rizzo, <i>The Little Network Book for Windows and Macintosh</i> (1999)
Appendix I	Excerpts from Alan Neibauer, <i>This Wired Home: The Microsoft Guide to Home Networking</i> (2000)
Appendix J	Steve Rigney, <i>Print Servers</i> (PC Magazine, January 19, 1999)
Appendix K	Excerpts from Dan Gookin, <i>PCs for Dummies</i> (6th Ed., 1998)
Appendix L	Apple Press Release: <i>Apple Introduces AirPort Wireless Networking</i> (July 21, 1999)

Appendix M	Frank J. Derfler Jr. and Les Freed, <i>Wireless LANs</i> (PC Magazine, April 18, 2000)
Appendix N	Intersil Corp., PRISM™ II Chip Set Overview (February 1999)
Appendix O	Apple Press Release: <i>Apple Introduces PowerBook G3</i> (May 6, 1998)
Appendix P	Apple, <i>Macintosh PowerBook G3 Series</i> brochure (May 1998)
Appendix Q	Apple Press Release: <i>Apple Unveils New iBook Line</i> (Sept. 13, 2000)
Appendix R	Apple, iBook web page: “Presentations” tab (archived Oct. 19, 2000)
Appendix S	Excerpts from Mary Ann Pike and Noel Estabrook, <i>Using FTP</i> (1995)
Appendix T	Excerpt from Preston Gralla, <i>How the Internet Works</i> (1998)
Appendix U	Excerpts from Ian S. Graham, <i>HTML Sourcebook</i> (1995)
Appendix V	Excerpts from J. Postel and J. Reynolds, RFC 959: <i>File Transfer Protocol (FTP)</i> (October 1985)
Appendix W	Excerpts from T. Berners-Lee et al., RFC 1945: <i>Hypertext Transfer Protocol - - HTTP/1.0</i> (May 1996)
Appendix X	Excerpts from Jill Ellsworth, Bill Barron, et al., <i>The Internet 1997 Unleashed</i> (1997)
Appendix Y	Excerpts from Aviel D. Rubin et al., <i>Web Security Sourcebook</i> (1997)
Appendix Z	Excerpts from David Pogue, <i>The iBook for Dummies</i> (2000)
Appendix AA	Cisco Systems, <i>Cisco Aironet 340 Series Client Adapters and Access Points</i> , Data Sheet (Feb. 2000)
Appendix BB	Excerpts from W. Richard Stevens, <i>TCP/IP Illustrated, Vol. 3</i> (1996)

Appendix CC	J. Franks et al., RFC 2069: <i>An Extension to HTTP : Digest Access Authentication</i> (Jan. 1997)
Appendix DD	Excerpts from Rogers Cadenhead, <i>How to Use the Internet: 2001 Edition</i> (2000)
Appendix EE	Excerpts from Chuck Musciano & Bill Kennedy, <i>HTML & XHTML</i> (4th ed., August 2000)
Appendix FF	Excerpts from Thomas J. Fallon, <i>The Internet Today</i> (2001)
Appendix GG	Excerpts from Brian Underdahl and Keith Underdahl, <i>Internet Bible</i> (2nd edition, 2000)
Appendix HH	Excerpts from Roy Hoffman, <i>Data Compression in Digital Systems</i> (1997)
Appendix II	Excerpts from Ronald K. Jurgen, <i>Digital Consumer Electronics Handbook</i> (1997)
Appendix JJ	Excerpts from <i>Bluetooth Core Specification v1.0 B</i> (December 1, 1999)
Appendix KK	Excerpts from David Pogue, <i>Mac OS 9: The Missing Manual</i> (2000)
Appendix LL	Christopher Breen, <i>Cut Loose</i> (discussing Apple's Airport Technology), MacWorld (June 2000)
Appendix MM	Erik Guttman, <i>Service Location Protocol: Automatic Discovery of IP Network Services</i> (IEEE Internet Computing, July / August 1999)
Appendix NN	Mark R. Brown, <i>Using Netscape Communicator 4</i> (1997)
Appendix OO	Yaron Goland et al., IETF Draft: <i>Simple Service Discovery Protocol/1.0</i> (Oct. 28, 1999)
Appendix PP	Excerpts from Cisco Systems, <i>Using the Cisco Aironet 340 Series Access Point</i> (2000)

Appendix QQ	Excerpts from Axis Communications, <i>Axis 540/640 Network Print Server User's Manual</i> (Sept. 1997)
Appendix RR	Excerpts from Abdelsalam (Sumi) Helal et al., <i>Any Time, Anywhere Computing</i> (1999)
Appendix SS	Thomas E. Truman et al., <i>The InfoPad Multimedia Terminal: A Portable Device for Wireless Information Access</i> (IEEE Transactions on Computers, Oct. 1998)
Appendix TT	Excerpts from Michael Miller, <i>The Complete Idiot's Guide to Home Theater Systems</i> (2000)
Appendix UU	Jaap Haartsen, <i>BLUETOOTH—The Universal Radio Interface for Ad Hoc, Wireless Connectivity</i> (Ericsson Review No. 3, 1998)
Appendix VV	Excerpts from Martin Doucette, <i>Digital Video for Dummies</i> (1999)
Appendix WW	Adobe Systems, <i>Adobe Premiere 5.0 At a Glance</i> (1998)
Appendix XX	Excerpts from Douglas W. Allen and Steve Johnson, <i>The Learning Guide to the Internet</i> (1997)

I. INTRODUCTION AND SUMMARY OF TESTIMONY

1. I, Samrat Bhattacharjee, have been retained by Petitioner Roku, Inc. (“Roku”) to investigate and opine on certain issues relating to United States Patent No. 11,029,903 (“the ’903 patent”) in Roku’s Petition for Post Grant Review of that patent. The Petition requests that the Patent Trial and Appeal Board (“PTAB” or “Board”) review and cancel claims 1-20 of the ’903 patent.

2. I am being compensated for my work on this matter by Roku for consulting services including time spent testifying at any hearing that may be held. I am also reimbursed for reasonable and customary expenses associated with my work in this case. I receive no other forms of compensation related to this case. My compensation does not depend on the outcome of this post grant review or the co-pending district court litigation, and I have no other financial interest in this post grant review.

3. This declaration is based on the information currently available to me. To the extent that additional information becomes available, I reserve the right to continue my investigation and study, which may include a review of documents and information that may be produced, as well as testimony from depositions that have not yet been taken.

4. I understand that the ’903 patent has been assigned to Flexiworld Technologies, Inc. (“Flexiworld” or “Patent Owner”).

A. Qualifications

5. My qualifications for forming the opinions in this expert report are summarized here and more fully detailed in my CV attached hereto as Appendix A.

6. I received Bachelor of Science degrees in both Computer Science and in Mathematics from Georgia College in 1994, and a Ph.D. in Computer Science in 1999 from Georgia Tech. My Ph.D. research was in developing a new form of networking architecture, and part of the work I did focused heavily on better delivery of video over the Internet. After receiving my Ph.D., I joined the University of Maryland as an Assistant Professor in 1999. In 2005, I was promoted to Associate Professor with tenure, and to Full Professor in 2009. At Maryland, I have taught both undergraduate and graduate courses in Computer Networking, Operating Systems, Computer Security, and various special topics courses on topics in related fields. My courses cover the basic structure of Computer systems and networking, and some cover media content delivery over the Internet in detail.

7. Both as a graduate student and as a faculty member, I have published in the top venues in Computer Networking, Computer Systems, and in Security. The list of my publications is attached as part of my CV in Appendix A. My research work has been supported by multiple grants from the US National Science

Foundation, and the Department of Defense. I have also started a Joint Ph.D. program with the University of Maryland and the Max Planck Society in Germany, and co-founded the annual Cornell, Maryland, Max Planck Research School that provides research exposure to about 80 students from across the world during a week-long school.

8. As I mentioned earlier, part of my Ph.D. research was to develop new architectures for video delivery on the Internet, and I have published papers on this architecture during my graduate studies. I continued to work on video delivery as a faculty member, and have published various papers on video streaming, content delivery architectures, and on resilient large-scale content delivery. During 2007, I was a visiting researcher at AT&T Labs, and one of the projects I focused on was a video content delivery platform. This work resulted in both publications and a granted US patent (U.S. Pat. No. 8,752,100 B2).

B. Materials Considered

9. Among the materials I reviewed in forming my opinions are the '903 patent, the prosecution history of the '903 patent, Exhibit 12 to Flexiworld's complaint which sets forth infringement allegations for the '903 patent, and the Exhibits and Appendices referenced in this declaration. I have also relied on my own professional and academic experience and my experience with working with others involved in the industry.

C. Level of Ordinary Skill in the Art

10. It is my opinion that a person of ordinary skill in the art (“POSA”) at the time of the invention would have had (1) a bachelor’s degree in computer science or computer engineering or a similar field, and (2) two years of experience developing software. The POSA would be familiar with well-known networking and web technologies. This description is approximate, in the sense that additional experience could make up for less education and vice versa.

11. I understand Flexiworld has not yet identified an alleged priority date for any claims of the ’903 patent in the district court litigation. In my view, the level of ordinary skill in the art would be similar regardless of whether the claims are entitled to a priority date as early as November 1, 2000 based on the earliest filed provisional application or if the claims are only entitled to a priority date of November 26, 2019 based on the filing of the ’903 patent’s actual application. Of course, a POSA in 2019 would have additional knowledge of newer technologies (*e.g.*, the iPhone), but none of the claims require technologies that would not have been known to a POSA on November 1, 2000.

D. Summary of opinions

12. Throughout my analysis and in forming all the opinions stated in this declaration, I have considered the perspective of a person of ordinary skill in the art at the time of the alleged invention.

13. It is my opinion that claims 1-20 of the '903 patent are invalid for lack of written description.

14. It is my opinion that claims 1-20 recite technology that was well-understood, routine, and conventional by late 2000 and even more well-understood, routine, and conventional by late 2019 when the '903 patent was filed.

II. OVERVIEW OF THE TECHNOLOGY

A. Relevant State of the Art

1. Wireless Communications

15. Although wireless networking technology had existed for years, standardization efforts in the late 1990s spurred increased interest in and use of wireless. The first major international wireless local area network (LAN) standard was IEEE 802.11. *See* Appx. B (Geier, 1999) at 89-96 (introducing the 802.11 standard). The initial 802.11 standard was finalized in 1997 and supplements in 1999 covered extensions (802.11a and b) that provided for increased data rates. IEEE 802.11 quickly came to dominate the wireless LAN space and replace earlier proprietary wireless technologies. By the time of the alleged invention (no earlier than November 1, 2000), IEEE 802.11 was essentially synonymous with wireless LAN technology. A New York Times article from October 30, 2000 describes surging enthusiasm around IEEE 802.11 wireless LAN technology. Appx. C (Markoff); *see id.* at 1 (“There is no doubt, however, that ‘wireless Ethernet’--

formally known as the 802.11b wireless technical standard as specified by the Institute of Electrical and Electronics Engineers -- is finally taking off.”).

16. Another important wireless standard, Bluetooth, was adopted in 1999. Bluetooth was developed to support low power radio connections between electronic devices, including computers and peripherals such as printers. *See generally* Appx. D (Haartsen, *The Bluetooth Radio System*, Feb. 2000); *see id.* at 6 (“The Bluetooth technology ... eliminates the need for wires, cables, and the corresponding connectors between cordless or mobile phones, modems, headsets, PDAs, computers, printers, projectors, and so on, and paves the way for new and completely different devices and applications.”).

17. Wireless networking is largely confined within the lower layers of the networking stack, *i.e.*, the physical layer, data link layer, and sometimes the data link layer.

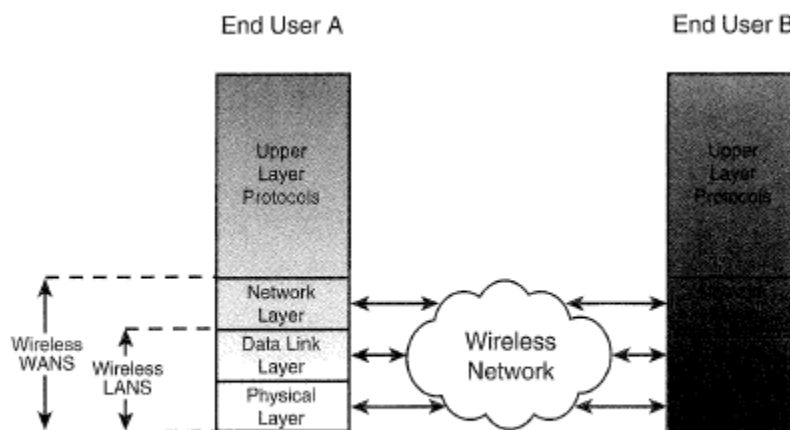


FIGURE 1.11 Wireless LANs and MANs fulfill Data Link and Physical Layer functionality; whereas, wireless WANs also include functions at the Network Layer.

Appx. B (Geier, *Wireless LANs*, 1999) at 39 (“As shown in Figure 1.11, wireless networks operate only within the bottom three layers [of the seven-layer OSI networking model].”). Transport layer protocols such as TCP “shield[] the higher layers from the networking implementation details.” *Id.* at 38. As a result, many upper layer / application processes that rely on networking (*e.g.*, network printing, accessing a network share, web browsing) are not significantly impacted by the use of wireless networking.

2. Service Discovery Technologies

18. In the late 1990s portable and handheld computing devices were becoming increasingly popular and it was generally understood that the utility of these devices could be enhanced by enabling them to discover and interact with other computing devices. To that end, a number of “service discovery technologies were developed ... to simplify the use of mobile devices in a network by allowing them to be ‘discovered,’ configured and used by other devices with a minimum of manual effort.” Appx. E (Richard, *Service Advertisement and Discovery: Enabling Universal Device Cooperation*, 2000) at 18; *see also* Appx. F (Networking Encyclopedia, 2001) at 1131-33 (“Service Advertising and Discovery”).

19. Universal Plug and Play (UPnP) is a technology platform developed by the UPnP Forum led by Microsoft. UPnP includes Simple Service Discovery

Protocol (SSDP) “for service discovery and advertisement.” Appx. E (Richard) at 23; *id.* at 24 (“In SSDP, each service has three associated IDs—service type, service name, and location—which are multicast when services are advertised.”). Apart from service discovery, UPnP includes a range of complementary technologies that facilitate interoperability between networked devices. *Id.* at 24 (discussing description, control, and presentation functionality).

20. In addition to UPnP, there were several other technologies that provided similar service discovery functionality. The Richard article discusses Jini, Salutation, and SLP, for example. Appx. E at 20-25. Bluetooth included a service discovery protocol (SDP) that “provides a simple API for enumerating the devices in range and browsing available services.” *Id.* at 19; *see also* Appx. F (Networking Encyclopedia, 2001) at 1131-33 (discussing Salutation, SLP, Microsoft.NET, SSDP, Bluetooth, Jini, JetSend, and Inferno).

3. Print Servers

21. The '903 patent states that an output controller for a printer can be a print server. '903 patent at 14:59-61 (“Other possible implementations of output controller 104 may include, for example, a ... print server.”); *see also id.* at 18:20-21; 19:60-64; 24:2-4. Because many printers did not include built-in network-interface cards (“NICs”), print servers could be used to connect printers lacking such cards to networks.

There are two basic methods for connecting your printers directly to the network. You can use a high-end printer that comes with a network card that is either built in or available as an option. Or you can use a stand-alone network print server—the Hewlett-Packard JetDirect EX is a good example—that supports a variety of protocols and usually comes with drivers to support many network operating systems, including Windows NT server.

Appx. G (Russel, *Running Windows NT Server 4.0*, 1997) at 220; *see also, e.g.*,

Appx. H (Poole, *The Little Network Book for Windows and Macintosh*, 1999) at

130 (“If the printer doesn’t have a built-in network port, you’ll have to install and configure the internal or external network adapter, or print server, that was described earlier in this chapter.”).

22. Some print servers were embodied as cards that were physically installed in the printer. *Id.* (“Installing an internal print server usually means inserting an adapter card into the printer’s expansion slot.”); Appx. I (Neibauer, *This Wired Home*, 2000) at 245 (“For some HP LaserJet printers, you can purchase an internal print server that fits inside the printer, much the way some NICs fit inside a computer.”). External print servers, on the other hand, connected to the printer by cable (e.g., parallel or USB cable). *Id.* at 247-249 (discussing setup of external print server). Appx. J is a PC Magazine feature from January 1999 comparing various external print servers including products from Axis, D-Link, HP, Intel, Lexmark and Linksys.

B. The '903 Patent

1. The '903 Patent Purports to Eliminate the Need to Pre-Install Drivers for Printers and Other Output Devices.

23. The '903 patent suggests that the invention relates to way of enabling users to output content from their information apparatuses to output devices without having to install device drivers before doing so. For example, the patent suggests “there is a need to provide in a manner in which a user can more conveniently or easily output digital content to an output device *without the inconvenience of finding and installing new device drivers or printer drivers.*” ’903 patent at 3:45-49.¹ The patent also suggests that “a more convenient or automated printing and output solution is needed so that a user can simply walk up to an output device (e.g., printer or display device) and easily output a digital document *without having to install or pre-install a particular output device driver (e.g., printer driver).*” *Id.* at 4:4-10.

24. The '903 patent critiques the “conventional[]” process for outputting content from an information apparatus to an output device (e.g., printer) is problematic because it requires users to install device drivers. *Id.* at 2:59-3:13; 13:44-47. “For a home or office user, this installation process may take anywhere

¹ All emphasis used when quoting source material in this declaration has been added unless otherwise noted.

from several minutes to several hours” *Id.* at 3:28-38. Device drivers are even more of a problem, according to the patent, for “mobile device users.” *Id.* at 3:50-55. The patent suggests it is not practical for mobile device users to pre-install drivers for all the output devices they may encounter while traveling. *Id.* at 3:55-64. Additionally, “many mobile information apparatuses have limited memory space, processing capacity, and power.” *Id.* at 4:11-13. Installing and running device drivers uses up that limited memory space. *Id.* at 4:16-19. Running device drivers on mobile devices can be slow due to their limited processing capacity and can drain their limited power. *Id.* at 4:20-25.

25. The ’903 patent purports to describe an invention that solves these problems and enables users to output content from their information apparatuses to output devices without the need to pre-install a device driver:

One implementation of the present invention provides an easy, friendly and convenient process for digital output. ***Unlike conventional output or printing, a user does not have to manually pre-install a device driver (e.g., printer driver)*** from a CD, floppy disk, or download the driver somewhere from a network. This is well-suited for providing output capability to small and lower-cost mobile devices with limited memory space, power supply and processing capability to still be able to output or print to an output device.

Id. at 4:48-56.

26. The way that the alleged invention avoids the need to install device drivers is by processing content for output not at the information apparatus itself

but instead at a “remote application server 110”:

In pervasive output operations of the present invention described below, various *device specific drivers* or applications may be *available* and may be *executed* completely or partially *in a remote application server 110*, thereby reducing the workload of information apparatus 100 and realizing device-independent pervasive output.

Id. at 13:51-57. Relatedly, the patent states that server application 112 may include “[c]omponents and operations to process the objects received to generate device-dependent output data acceptable to one or more output devices 106 selected by a user.” *Id.* at 18:4-7; *see also id.* at 21:58-61 (“A server application 112 obtaining and processing the document object and converting it into output data, reflecting at least in part a relationship to said output device object;”); 22:61-23:2 (“The processing and generation of output data [by server application 112] may reflect at least in part a relationship to the output device object and or job object contained in the composite message received from client application 102.”); 29:45-32:35 (describing server application process in connection with Fig. 7).

2. The ’903 Patent’s Pervasive Output Process

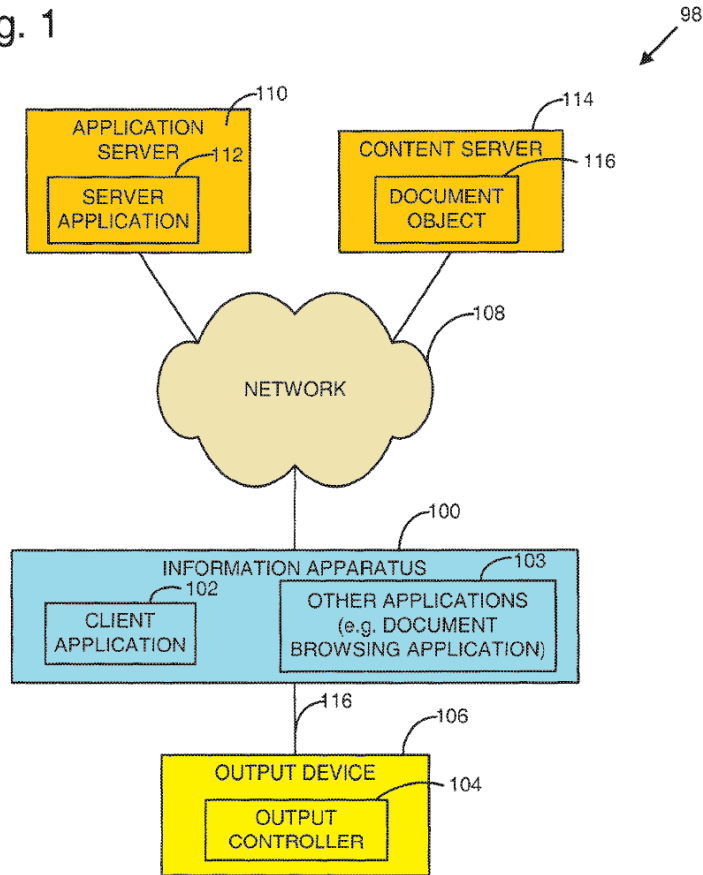
27. The ’903 patent states that the “[p]resent invention relates to providing digital content to an output device and, in particular, to providing pervasive output in which an *information apparatus can pervasively output digital content to an output device regardless of the processing power, display*

screen size and memory space of the information apparatus.” ’903 patent at 1:49-56.

28. This “pervasive output” process involves processing content for output at a remote application server rather than at the user’s device, as just discussed. *Id.* at 13:51-57 (“*In pervasive output operations of the present invention* described below, various *device specific drivers or applications* may be *available* and may be *executed* completely or partially *in a remote application server 110*, thereby *reducing the workload of information apparatus 100 and realizing device-independent pervasive output.*”). Because the pervasive output process does not require the information apparatus to display or process content for output devices *itself*, the process purportedly works “regardless of the processing power, display screen size and memory space of the information apparatus.”

29. Figure 1 (below, annotations added) “is a block diagram of a *pervasive output system 98* that can implement the process and apparatus of present invention.” *Id.* at 8:59-61.

Fig. 1



Information apparatus 100 communicates with application server 110 and content server 114 over network 108. *Id.* at 8:61-65. Network 108 may include a wide area network (WAN) or the Internet. *Id.* at 9:3-9.

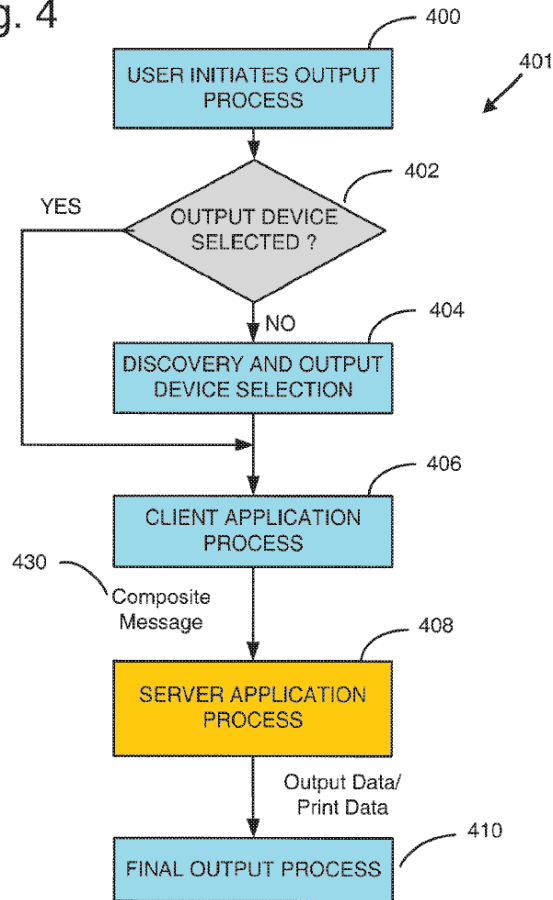
30. “Information apparatus 100 is a computing device with processing capability.” *Id.* at 9:10-11. “[I]nformation apparatus 100 may be a mobile computing device such as palmtop computer, handheld device, laptop computer, personal digital assistant (PDA), smart phone, screen phone, e-book, Internet pad, communication pad, Internet appliance, pager, digital camera, etc.” *Id.* at 9:11-18. “A typical example of output device 106 may be a printer” *Id.* at 11:65-66; *see*

also id. at 12:1-20 (identifying other possible output devices including “monitors,” “projectors,” and “sound output devices”).

31. Information apparatus 100 communicates with an output device 106 over communication link 116. *Id.* at 8:65-67. Communication link 116 “may be a ***short-range radio interface*** such as those implemented according to the Bluetooth or IEEE 802.11 standard.” *Id.* at 9:57-61. Thus, in contrast to network 108, there is no suggestion that communication link 116 involves communication over a WAN or the Internet.

32. The patent suggests that an application on information apparatus, namely “***pervasive output client application 102***,” is what “provides [the] ***pervasive output capability of the present invention.***” *Id.* at 10:56-60; Fig. 1. The “functionalities and process of pervasive output client application 102 are described ... in the pervasive output process with reference to FIG. 4.” *Id.* at 11:59-61. “FIG. 4 is a flow diagram of a ***pervasive output process 401 of the present invention.***” *Id.* at 21:45-46.

Fig. 4



Id. at Fig. 4 (annotations added). As I noted above, the patent states that this “[p]ervasive output process 401 allows an information apparatus 100 to output digital content or document in its original form to an output device 106 regardless of processing power, display screen size, or memory space of information apparatus 100.” *Id.* at 21:46-50.

33. “Typically, a user initiates output process 401 by invoking a client application 102 in his/her information apparatus 100.” *Id.* at 22:4-6. “[A] user may need to select one or more output devices 106 for output service.” *Id.* at 22:19-20. “An optional discovery process 404 may be implemented to help the

user select an output device 106.” *Id.* at 22:19-22. The patent identifies known “standards or protocols” such as Bluetooth and Universal Plug and Play that can be used to implement the optional discovery process. *Id.* at 23:66-24:9.

34. “In stage 406, the client application 102 ... obtain[s] (1) a document object, (2) an output device object and (3) any other optional objects such as a job object.” *Id.* at 22:42-46. “The client application 102 may create a composite message including these objects ... and transmit the composite message to the server application 112 for processing” *Id.* at 22:56-60.

35. “The server application 112, after receiving such a composite message from the client application 102, may in step 408 processes [sic] the document object or objects contained in the composite message and convert it or them into output data.” *Id.* at 22:61-65. “The processing and generation of this output data may reflect at least in part a relationship to the output device object and or job object contained in the composite message received from client application 102.” *Id.* at 22:66-23:2. “The output data generated may be transmitted back to the information apparatus 100, requesting output service of process 401 via network 108.” *Id.* at 23:3-5. “In step 410, information apparatus 100 transmits output data, with or without further processing, to the selected output device 106 through a local communication link 116.” *Id.* at 23:8-10.

36. The ’903 patent thus describes *client application 102 on information*

apparatus 100 sending certain data (e.g., “document objects,” “output device objects,” and, optionally, “job objects”) to server application 112, receiving processed “output data” from the server application, and transmitting the output data to output device 106. The ’903 patent further describes these steps in connection with Figs. 5-8 which correspond the optional discovery process, the client application process, the server application process, and the final output process, respectively. *Id.* at 5:24-32 (brief description); *id.* at 23:46-33:27 (detailed description).

37. Finally, all the devices discussed in the ’903 patent (*i.e.*, the information apparatus, the output controller, the output device, and the application server) are described in functional terms. The ’903 patent specification lists eleven “functionalities” that may be included in client application 102 (*id.* at 11:5-61); eleven functionalities that may be included in output controller 104 (*id.* at 15:34-16:25), and twelve functionalities that may be included in output controller 106 (*id.* at 16:52-40), and ten functionalities that may be included in server application 112 (*id.* at 17:65-18:57). In some cases, these functionalities are referenced as “components” and/or “operations” but the components and operations are only identified and discussed based on the function they perform. *See, e.g., id.* at 15:36-38 (output controller 104 may contain “[i]nput components and operations for receiving service requests from a plurality of information apparatuses 100”).

C. Claim Construction

38. I understand that in a PGR proceeding, the challenged claims are construed “in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.” In evaluating the challenged claims, I have applied my understanding as to what a person of ordinary skill in the art would have understood these claims to mean as of November 1, 2000 when the earliest-filed priority application was filed.

III. UNPATENTABILITY OF THE CHALLENGED PATENT CLAIMS

A. Relevant Legal Standards

1. Lack of Written Description Under 35 U.S.C. § 112

39. I understand that that there is a written description requirement that requires that a patent specification “shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same.” I have been advised that the test for sufficiency of the written description is whether the disclosure in the specification reasonably conveys to those skilled in the art that the inventors had “possession” of the claimed subject matter as of the application’s filing date. I understand that “possession” of the claimed subject matter must be demonstrated by the material within the four corners of the specification. In other

words, the specification must describe an invention understandable to a person of ordinary skill in the art and show that the inventor actually invented what is claimed.

2. Subject Matter Eligibility Under 35 U.S.C. § 101

40. I understand that where a patent claim is directed to an abstract idea, the claim is invalid unless the claim contains an “inventive concept” which must be an element or some combination of elements that ensures the patent claim amounts to significantly more than a patent claim on the abstract idea itself. I understand that an inventive concept cannot be well-understood, routine, conventional activities previously known to the industry. I also understand that a “wholly generic computer implementation” is not generally sufficient to provide an inventive concept. I understand that the elements of a claim must be considered both individually and also as an ordered combination in determining whether they include an inventive concept.

B. Lack of Written Description

41. As I discussed above, the '903 patent describes the “pervasive output client application 102” on information apparatus 100 that sends certain data (e.g., certain “objects”) to server application 112, receives processed “output data” from the server application, and then transmits the output data to output device 106. *See supra* ¶¶ 27-36.

42. The challenged claims describe a substantially different system / approach in which many of the functions described as being performed by the “pervasive output client application 102” on information apparatus 100 are performed instead by an “output system” (claims 1 and 15) or a “wireless output device” (claim 8). The claimed system is not supported by the written description, in my opinion.

1. The Written Description Fails to Support the Independent Claims.

a. The '903 Patent's Written Description Does Not Support an Output System / Device that Connects to Servers Over the Internet.

43. Claim 1, element 1[h] recites that the output system connects to one or more servers over the Internet:

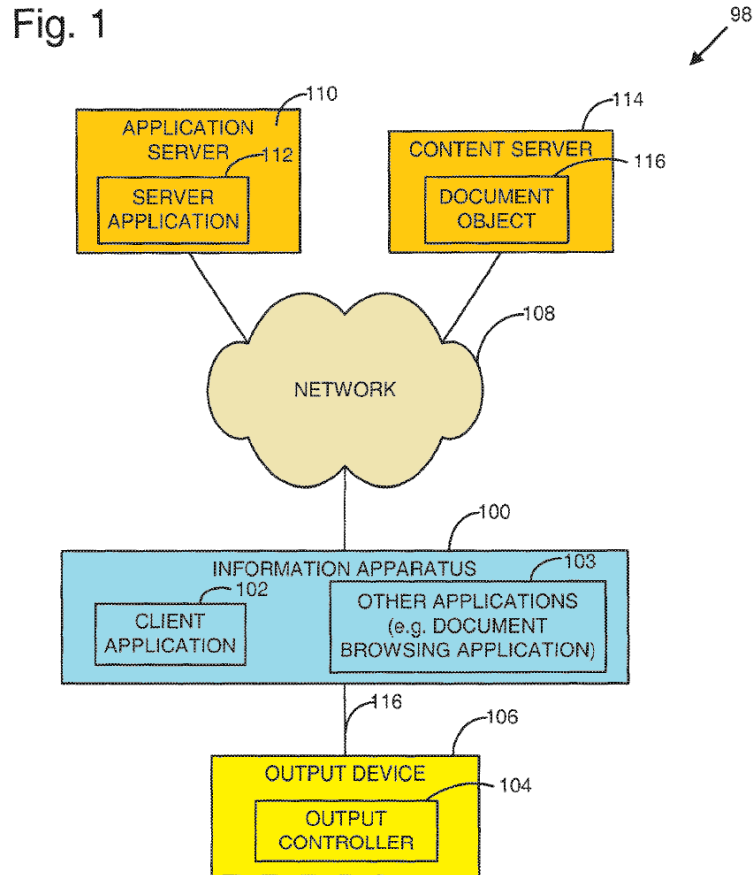
(3) wirelessly connecting the output system, by the output system, using the at least one chip or chipset of the output system, and over the wireless local area network wirelessly coupled in (2), to the one or more servers over the Internet;

Ex. 1003. Element 8[f] of claim 8 and element 15[f] of claim 15 are very similar.

Id.

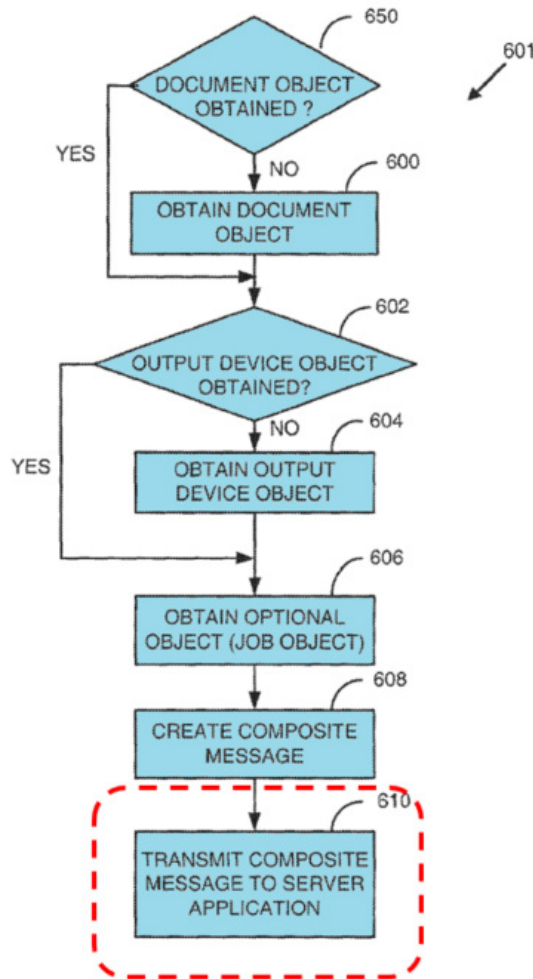
44. In my opinion, these limitations are not supported by the '903 patent's written description. The patent does not describe any *output system* connecting to any server over the Internet. Instead, the patent describes the information apparatus connecting to servers over the Internet.

45. The patent describes two servers: application server 110 and content server 114. Both servers are shown in Fig. 1 (below, annotations added):



46. First, with respect to application server 110, the specification describes *client application 102* connecting to application server 110 in order to send a composite message. This is depicted, for example, with step 610 of Fig. 6 which depicts “exemplary client application process 601”:

Fig. 6



'903 patent at 25:23-36 (“[E]xemplary client application process 601 ... may include or utilize ... [a] client application 102 that coordinates with server application 112 to manage the process of communication and transmission of a composite message (including at least a printer object and a document object) to the server application 112 for further processing.”); *see also id.* at 22:56-60 (“The client application 102 may create a composite message including these objects (document object, output device object and other optional objects) and transmit the composite message to server application 112 for processing, as described below in

greater detail with reference to FIG. 6.”). In describing step 610, the specification notes that “client application 102 may communicate with the server application 112 using one or more or a combination of standard network protocols” and that “*communication link between information apparatus 100 and application server 112* may be implemented with one or a combination of standard *network connections and communication links.*” *Id.* at 28:55-29:2.

47. Second, with respect to content server 114, the specification again identifies information apparatus 100 as the device that connects to content server 114. The specification states, for example: “The digital documents stored in content server 114 may be viewed or edited by a user *using an information apparatus 100.*” *Id.* at 19:3-5. The specification goes on to describe using a web browser or other application on information apparatus 100 to access the content server. *Id.* at 19:5-20. Consistent with this discussion of content server 114, the patent’s description of information apparatus 100 notes that it may include document browsing application 103 for viewing digital documents stored “in a network node (e.g., in content server 114).” *Id.* at 10:18-39. Nowhere in the patent is there any description of any document browsing application or analogous software on output device 106 or output controller 104.

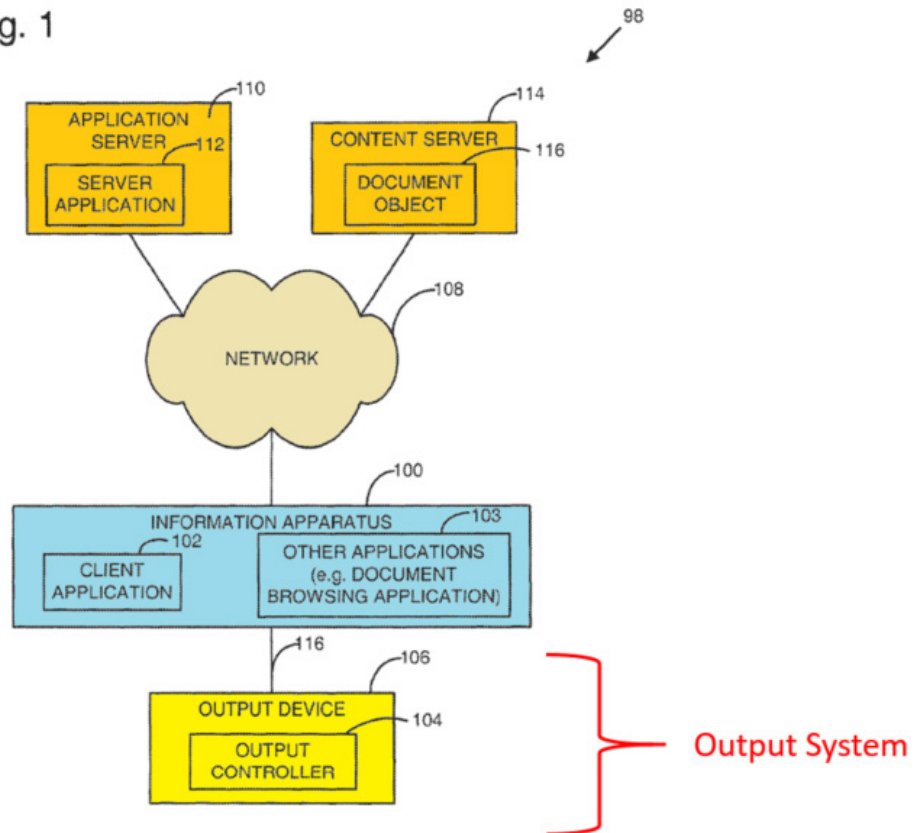
48. The patent’s description of the “functionalities” of the various components is also consistent. The specification states that client application 102

may include functionality to “[c]oordinate with a server application 112 residing in application server 110 to *manage the process of communication and transmission of objects or data to and from application server 112.*” ’903 patent at 11:14-17.

Relatedly, server application 112 may include “[c]omponents and operations to manage and coordinate *communication with an information apparatus 100* requesting output service.” *Id.* at 18:31-33. None of the server application 112 functionalities reference communication with output device 106, output controller 104, or any output system. *See id.* at 17:65-18:57. Nor do output device 106 or output controller 104 include functionalities for connecting to or communicating with any server. *See id.* at 15:34-16:25 (output controller 104 functionalities), 16:53-17:40 (output device 106 functionalities).

49. Finally, all of this is consistent with Figure 1 (below, annotations added) which shows that while information apparatus 100, application server 110, and content server 114 are all directedly connected to the same network 108, *output device 106 is not.*

Fig. 1



The only communication link from output device 106 shown in the figure is communication link 116 to information apparatus 100. The '903 patent asserts that a benefit of the invention is that the output device need not have a static network connection:

Finally, one implementation provides a convenient method allowing users to output to an output device with or without connection to a static network. Through local communication and synchronization between information apparatus and output device, hardware and software installation for static or permanent network connectivity may not be necessary for the output device.

Id. at 4:65-5:4. The patent also explains, “[f]or example, information apparatus 100 illustrated in FIG. 1 may communicate with output device 106 through a

Bluetooth standard interface while communicating with other network nodes (e.g., content server 114 or application server 110) through a cellular telephone modem interface.” *Id.* at 9:38-43.

b. The '903 Patent's Written Description Does Not Support an Output System / Device that Sends the Job Object and Document / Content Object to the Server.

50. Claim 1, elements 1[i] and 1[l], respectively, recite that the output system sends a job object and digital document object to the server:

(4) wirelessly sending, by the output system, using the at least one chip or chipset of the output system, and over the wireless local area network wirelessly coupled in (2), a job object ... from the output system to at least one server of the one or more servers over the Internet ...

...

(7) wirelessly sending, by the output system and using the at least one chip or chipset of the output system, from the output system, over the wireless local area network wirelessly coupled in (2), a digital document object to at least one server of the one or more servers over the Internet ...

Ex. 1003 (claim listing). Elements 8[g] and 8[k] in claim 8 and elements 15[h] and 15[k] in claim 15 recite similar limitations. *Id.*

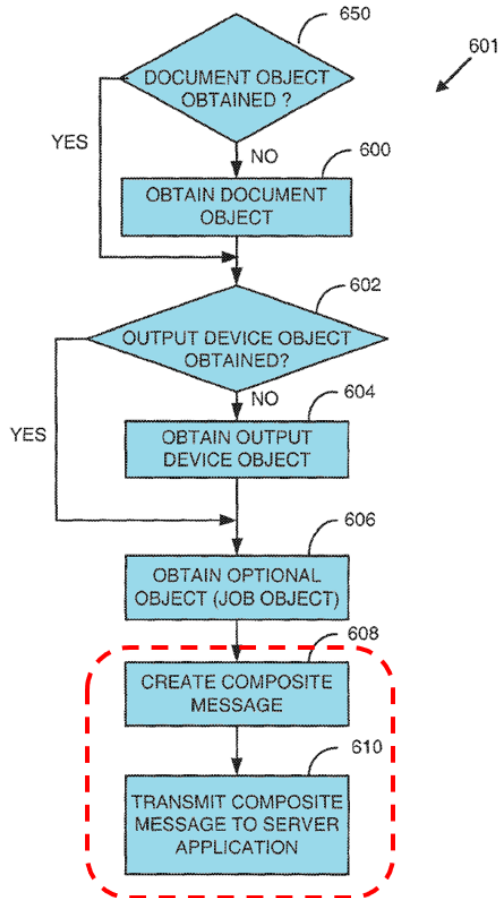
51. In my opinion, these limitations are not supported by the '903 patent's written description. Nowhere in the written description is there any statement or suggestion that the *output system* sends job or content objects to the server application. The written description makes clear that it is the user's information

apparatus 100, which contains the pervasive output client application 102, that sends the job and document objects to the server.

52. The written description explains that “The *client application 102 may create a composite message including these objects* (document object, output device object and other optional objects) and *transmit the composite message to server application 112* for processing, as described below in greater detail with reference to FIG. 6.” *Id.* at 22:56-60 (describing step 406 of Fig. 4); *see also id.* at 21:52-57 (“Pervasive output process 401 may include or utilize ... A *client application 102 transmitting objects to a server application 112*;”). Nothing in the '903 patent suggests that any part of client application 102 runs on the output device 106 or output controller 104, or that either of those output devices includes any equivalent application. The idea of relying on the output system to send these objects to the server would also be incongruous with the disclosures I discussed above about the output system not needing to have static or permanent network connectivity and this being a benefit of the alleged invention.

53. Figure 6 of the '903 patent (below, annotations added) shows “exemplary client application process 601” which includes transmission of the composite message with the various “objects” to server application 112. *Id.* at 25:23-36.

Fig. 6



In describing Figure 6, the patent states:

In step 608, the client application 102 may create or assemble a composite message. A composite message may be any type of data transferred across network 108 that may include one or more transmissions. *A composite message typically includes partially or entirely the objects (with some default values) obtained by the client application 102 in previous steps.*

In step 610, the client application 102 transmits the composite message to server application 112.

Id. at 28:48-56; *see also e.g., id.* at 25:23-36 (describing client application process 601 as involving client application 102 obtaining “objects” and transmitting them to server application 112); 11:14-17 (“client application 102 may include ...

functionalit[y] ... [to c]oordinate with a server application 112 residing in application server 110 to manage the process of ***communication and transmission of objects or data to and from application server 112.***”).

54. The patent’s description of server application 112 is consistent and states that it receives the recited objects from ***the client application 102***. For example, Fig. 7 depicts “exemplary server application process 701” includes step 700, in which “server application 112 receives a composite message ***from client application 102***. As mentioned earlier, the composite message may include one or more of a ***document object***, an output device object, and an optional ***job object***.” *Id.* at 29:66-30:3; *see also id.* at 18:1-3 (server application 112 may include “[c]omponents and operations to ***receive data and/or objects*** (with at least a output device object and a document object) ***from client application 102.***”); 29:54-56 (“A server application 112 that ***receives a composite message*** (including an output device object and a document object) ***from client application 102.***”). The specification also describes a scenario where the server application 112 receives no job object initially and therefore “launch[es] a GUI ***in the information apparatus 100*** as described above ***to obtain*** partially or entirely ***the job object*** fields (job preference) from the user.” *Id.* at 32:8-16; Fig. 7 (step 709: “distribute a GUI in user’s information apparatus (optional)”).

c. The '903 Patent's Written Description Does Not Support an Output System / Device that Receives Indications of Content Selections From the User.

55. Claim 1, element 1[k] recites that the *output system* receives the output data from the server:

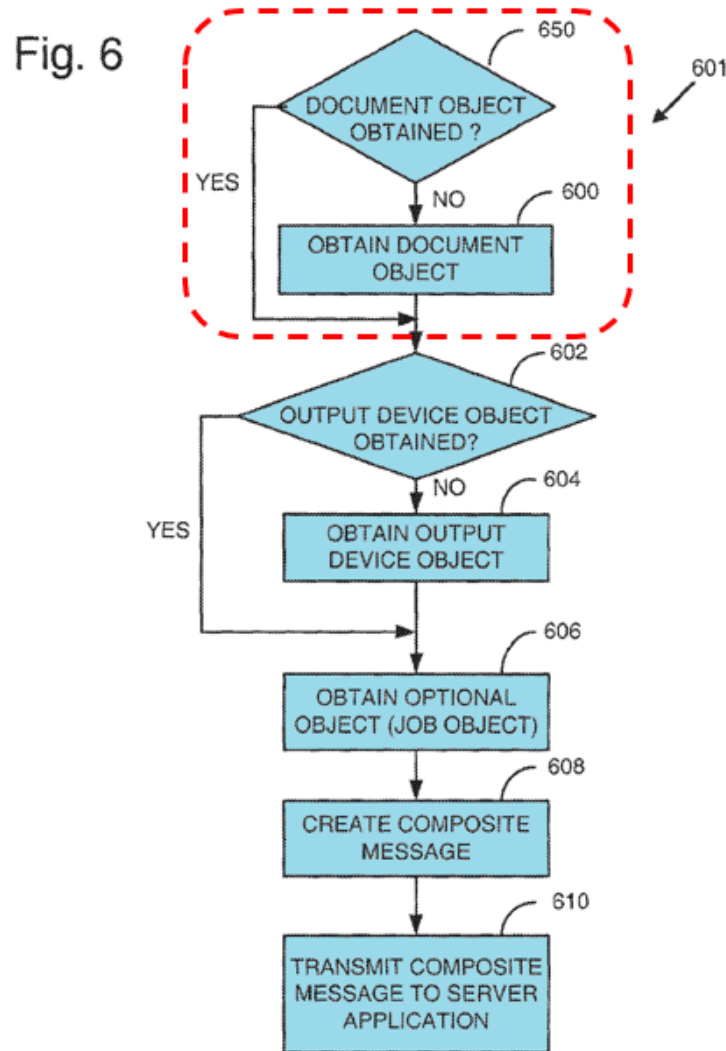
(6) *receiving, by the output system* and via the interface of the output system, an *indication of selected digital content* from among the one or more digital content that are available at the one or more servers;

Ex. 1003. Elements 8[j] in claim 8 and 15[j] in claim 15 recite very similar limitations. *Id.*

56. In my opinion, these limitations are not supported by the '903 patent's written description. The written description never states that the *output system* receives an indication from the user of selected digital content. Instead, it repeatedly and solely states that it is client application 102 on the user's information apparatus 100 that receives the indication of selected digital content. The second of the listed "functionalities" of client application 102 is functionality to "[o]btain document object (1) *from user input or selection*, or (2) from *other applications* (e.g., a document browsing application) residing *in the information apparatus 100*." *Id.* at 11:10-13. No similar functionality is identified for either output device 106 or output controller 104. *See id.* at 15:34-16:25 (output controller); 16:52-17:40 (output device).

57. The claimed step relates to step 600 of the "exemplary *client*

application process 601” shown in Figure 6 (below, excerpted and coloring added):



“Step 600 indicates that *client application 102* obtains a document object.” *Id.* at 25:37-38. “There are various ways that a *client application 102* may obtain a digital document or a pointer or reference to the digital document.” *Id.* at 25:49-51. The process may involve “a GUI provided by the client application 102”:

In one embodiment, the *client application 102* may provide a GUI with which a user can directly input the

pointer or reference (e.g., URL, IP address, filename, path, etc.) of a digital document stored locally or in a network node. This manual process may be facilitated by, for example, providing a GUI with which users may select one or more pointers from a list of pointers or references of digital documents stored locally or in a remote network node. For instance, ***through a GUI provided by the client application 102***, a user may see and select pointers of digital documents stored in a remote file server.

Id. at 25:51-61. Client application 102 runs on information apparatus 100. *Id.* at 10:57-60 (“[I]nformation apparatus 100 includes a pervasive output client application 102 that provides pervasive output capability of the present invention.”).

58. Alternatively, another application running on information apparatus 100 may indicate to client application 102 the content the user wants to input:

In another embodiment, the client application 102 may obtain output content or pointer to output content from ***another application in the same information apparatus 100***. As an example, a user may (1) launch the client application 102, and (2) invoke another application 103 (***e.g., document editing and or browsing application***) residing in the same information apparatus 100 to view or download the digital document. As another example, a user may (1) run another application 103 (e.g., document editing and or browsing application) residing in the same information apparatus 100 to view or download the digital document; and (2) launch or invoke the client application 102. In these cases, the client application 102 may communicate with another application 103 (e.g., document browsing application) to obtain pointers to the digital document and or the digital document itself (if it

has been downloaded locally for viewing) to be included in document object.

Id. at 25:64-26:14. This second approach is analogous to the conventional process by which users can print documents from other applications such as Microsoft Word (a “document editing” application) or a web browser (a “browsing application”).

59. Both of these approaches that identify a digital document for output involve a user interacting with the information apparatus 100 and its applications (client application 102 and/or other application 103). Neither approach references output device 106 or output controller 104.

60. Similarly, when the patent discusses content server 114, it states that the documents stored there are accessed using *information apparatus 100*, not output device 106:

Content server 114 may represent one of a plurality of server nodes on network 108 that may store digital documents 116. The digital documents stored in content server 114 may be viewed or edited by a user *using an information apparatus 100*. As an example, the content server 114 may be a web server that hosts a plurality of web pages written in mark up languages such as HTML, WML, XML, HDML, CHTML, among others. A user may view web pages using an Internet browsing application such as Internet Explorer or Netscape Navigator, a WAP browser, etc. As another example, the content server 114 may be a file server that allows multiple clients to store and share digital files with appropriate security or authentication procedures. These digital files or documents may contain one or more of

image, text, graphics, sound and video. The files may be saved in various file formats (e.g., MS Word, Excel, PowerPoint, PDF, Postscript, JPEG, GIF, MPEG, etc.). A user may need to have *appropriate application on his/her information apparatus 100* to access, view and edit these files.

Id. at 19:1-20. Elsewhere, the written description mentions that information apparatus 100 may document browsing applications such as a web browser. *Id.* at 10:18-39. Nowhere in the written description is there any suggestion that output device 106 or output controller 104 includes web browsers or other applications that would be capable of accessing files on content server 114.

d. The '903 Patent's Written Description Does Not Support an Output System / Device that Receives Output Data From the Server.

61. Claim 1, element 1[m] recites that the *output system* receives the output data from the server:

(8) wirelessly *receiving, by the output system*, using the at least one chip or chipset of the output system and over the wireless local area network wirelessly coupled in (2), *output data from at least one server* of the one or more servers over the Internet ...

Ex. 1003. Elements 8[l] of claim 8 and 15[l] of claim 15 recite very similar limitations. *Id.*

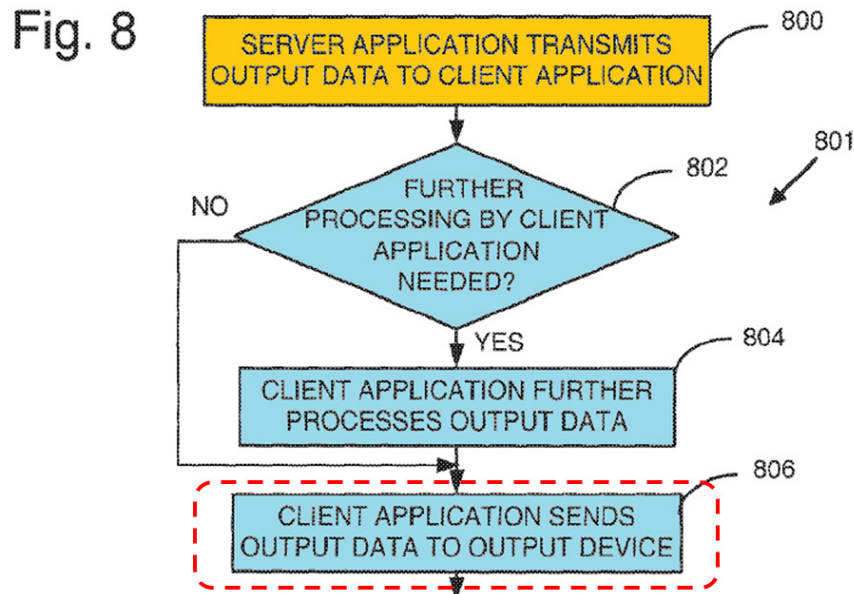
62. In my opinion, these limitations are not supported by the '903 patent's written description. The written description never states or suggests that the output system receives output data from the server. The written description repeatedly

describes *client application 102 on information apparatus 100* receiving the output data from the server. For example, the patent explains that server application 112 may contain “[c]omponents and operations to *transmit output data back to the client application 102.*” *Id.* at 18:10-11. None of the identified “functionalities” of the server application 112 reference sending output data to an output device or system. *Id.* at 17:65-18:57.

63. In describing the pervasive output process in Figure 4, the patent states that it may include “a server application 112 *transmitting output data to the information apparatus 100 including a client application 102.*” *Id.* at 21:62-64; *see also id.* at 23:3-5 (“The *output data* generated may be *transmitted back to the information apparatus 100*, requesting output service or process 401 via network 108.”). Similarly, in describing Figure 7, which depicts “an exemplary server application process 701,” the patent states that “[i]n step 710, the server application 112 *transmits output data to the information apparatus 100* through network 108.” *Id.* at 32:30-31.

64. Consistent with the above statements, the patent describes the *information apparatus 100* transmitting the output data to the output device 106. *See, e.g., id.* at 23:8-10 (“In step 410, *information apparatus 100 transmits output data*, with or without further processing, *to the selected output device 106* through a local communication link 116.”). This is also shown in Figure 8 which depicts

“final output process 801.” *Id.* at 32:36-45.



Id. at Fig. 8 (excepted, annotations added); *see also id.* at 32:46-48 (“After receiving output data from server application 112, the *information apparatus 100* may then transmit the output data to the output device 106 selected by the user in step 806.”); 32:41-43 (“[F]inal output process 801 for pervasive output ... may include or utilize ... An *information apparatus 100* transmits output data, with or without further processing, to the output device 106 selected by the user.”).

65. One benefit of this approach whereby data is downloaded to the information apparatus first and then transmitted from the information apparatus to the output device is that the output device need not necessarily be on the same network as the servers. As I discussed above, the ’903 patent suggests that a benefit of the invention is that the output device need not have a “static or permanent network connectivity.” *Id.* at 4:65-5:4; *see also id.* at 9:38-43 (giving

example where information apparatus communicates with servers via cellular network and output device 106 using Bluetooth).

66. The approach described in the '903 patent would in theory allow the “information worker at an airport receiving Email in his hand-held computer [to] walk up to a nearby printer or fax machine to have his e-mail printed” (*id.* at 2:30-32), without that printer or fax machine needing to be on any static network (*i.e.*, the same network as the servers that provide and/or process the document for output).

2. The Written Description Fails to Support the Dependent Claims for the Same Reasons.

67. The dependent claims of the '903 patent lack written description for the same reasons as the independent claims discussed above. None of the dependent claims narrow the independent claims to subject matter that is described. Dependent claims 3, 14, and 18 are briefly discussed below because they demonstrate one of the ways that the broader independent claims from which they depend include scope that is plainly not supported.


a. Dependent Claims 3, 14, and 18

68. Dependent claims 3, 14, and 18 depend from claims 1, 8, and 15 and recite that the “output system” (claims 3 and 18) or “wireless output device” (claim 8) is a “controller box” wire connected to a television. Ex. 1003. The “controller box” refers to an output controller implemented as a “box” connected to an output

device (in this case, a television). *See, e.g.*, '903 patent at 21:18-20 (“In the third implementation shown in FIG. 3C, the output controller 104C may be implemented in a separate **box** or server or station connected externally to output device 106C.”); Fig. 3C.

69. Dependent claims 3, 14, and 18 lack written description support based on the same limitations I discussed above for the three independent claims. Although the '903 patent discusses output controllers, there is no description of any output controller (a) connecting to any servers over the Internet; (b) sending job objects or document / content objects to any server; (c) receiving user content selections, or (d) receiving output data from a server.

70. Dependent claims 3, 14, and 18 also help to demonstrate that the broader independent claims are not described. In particular, these dependent claims confirm that the “output system” of claims 1 and 15 and the “wireless output device” of claim 8 extends to an external output controller (or “controller box”) even though the written description fails to describe any output controller performing the recited functions discussed above. *See, e.g.*, Ex. 1003, claim 3 (“The method according to claim 1, wherein the at least an output device is at least a television; and wherein ***the output system is embodied, at least in part, as a controller box wire connected to the television***”). The figure below illustrates this relationship.



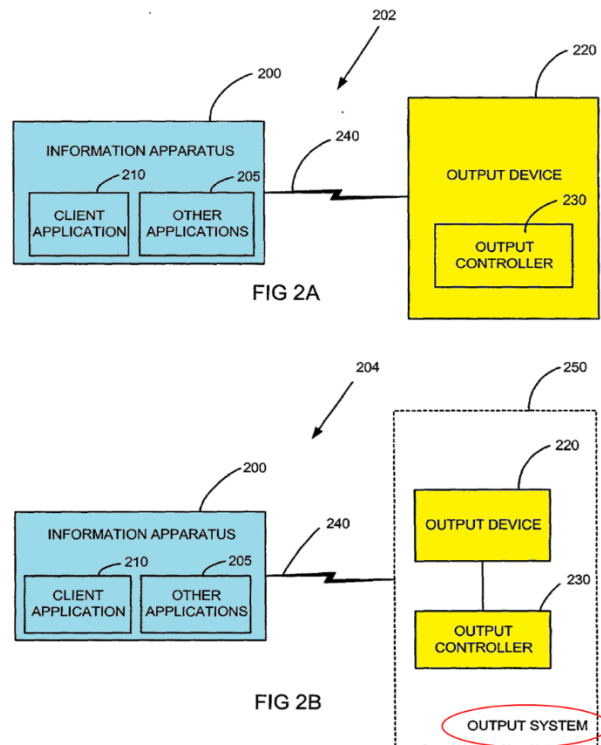
3. Information Apparatus 100 is not an “Output System” (claims 1 & 15) or “Wireless Output Device” (claim 8).

71. In my opinion, the information apparatus 100 that is described in the '903 patent would not itself be viewed by a POSA as an “output system” or “wireless output device” as those terms are used in the claims.

72. The term “output system” (claims 1 and 15) is used in the '903 patent to refer to systems consisting of an output device and an external output controller.² For example, some of the applications incorporated by reference into the '903 patent show and describe output system 250 consisting of an output

² The '903 patent also refers to the broader system, including information apparatus 100, network 108, and servers 110 and 114 as “pervasive output system 98” ('903 patent at 8:59-61), but this is clearly a different meaning than the claimed “output system” which includes device components (*e.g.*, a processor, a wireless communication unit, a user interface) but not the separately claimed “client device” or “one or more servers.”

device and an external output controller.



Ex. 1006 (U.S. Appl. No. 10/053,765) at Figs. 2A, 2B (color annotations added);

¶ [0076] (“The output system 250 includes an output device 220 and an output controller 230 which may be externally connected to, or otherwise associated with, the output device 220 in the output system 250.”). Note that Fig. 2B also shows an information apparatus 200 but it is not part of output system 250. The ’903 patent similarly shows in Fig. 3C a similar configuration involving an output device and an external output controller though it does not expressly refer to that configuration as an “output system.” ’903 patent at 21:18-29 (describing Fig. 3C). The term “wireless output device” (claim 8) is only used in the claims of the ’903 patent but suggests an output device that supports wireless communications.

73. The '903 patent clearly distinguishes between an information apparatus and an output device. The information apparatus is described as the user device that outputs (*i.e.*, transmits) processed output data to an output device that outputs the content (e.g., on paper, on a display, with sound) to the user. If the information apparatus is an adequate output system itself, the alleged invention makes little sense.

74. The first sentence of the patent specification explains that the invention relates to “providing pervasive output in which an ***information apparatus*** can pervasively output digital content to an ***output device*** regardless of the process power, display screen size and memory space of the information apparatus.” ’903 patent at 1:51-56; *see also id.* at 21:46-50 (“Pervasive output process 401 allows an ***information apparatus 100*** to ***output digital content or document*** in its original form ***to an output device 106*** regardless of processing power, display screen size, or memory space of information apparatus 100.”). A POSA would have understood that the “display screen size” of the information apparatus is not limiting because the content is output by a separate output device.

75. The patent asserts that “there is a need to allow users to easily output content and information ***from their pervasive computing devices to any output device.***” *Id.* at 2:23-25. The patent suggests that users want to “walk up to”

printers and other output devices and be able to easily output to them.³ The patent repeatedly discusses printers as the quintessential output devices, and none of the exemplary information apparatuses are devices that would have been able to print documents themselves.

76. Additionally, as I discussed, the patent explains the purported need for the invention in connection with problems associated with installing and running output device drivers on information apparatuses. But if an information apparatus were sufficiently capable to output content themselves, there would be no need to install any output device drivers to begin with, and no need for the alleged invention.

77. As I mentioned above, the patent lists different sets of “functionalities” for client application 102 of information apparatus 100, output device 106, and output controller 104. *Id.* at 11:5-58 (client application 102); 15:34-16:25 (output controller 104); 16:52-17:40 (output device 106). One of the listed “functionalities” of output device 106 is “[c]omponents and operations to

³ *Id.* at 2:30-32 (“To illustrate, an information worker at an airport receiving Email in his hand-held computer may want to **walk up to a nearby printer or fax machine** to have his e-mail printed.”); *id.* at 2:45-53 (“In still another example, a user with a mobile device may want to simply **walk up to a printer and conveniently print a file** that is stored on the mobile device or that is stored on a network (e.g., Internet, corporate network) and accessible from the mobile device, such as a PowerPoint® display application document, word processing document, or a document in any other file format such as PDF, HTML, JPEG etc.”).

output final output data or print data on a substrate or in another medium such as a display screen.” *Id.* at 16:63-65. No similar function is listed for client application 102 on information apparatus 100. *See id.* at 11:5-58. The patent also identifies different types of devices as examples of information apparatuses and output devices:

Examples of such information apparatuses include, without limitation, desktop computers, laptop computers, networked computers, palmtop computers (hand-held computers), personal digital assistants (PDAs), Internet enabled mobile phones, smart phones, pagers, digital capturing devices (e.g., digital cameras and video cameras), Internet appliances, e-books, information pads, and digital or web pads. ***An output device may include*** any one or more of fax machines, printers, copiers, image and/or video display devices (e.g., televisions, monitors and projectors), and audio output devices.

Id. at 1:61-2:4; *see also id.* at 9:10-18 (listing exemplary devices for information apparatus 100). 11:65-12:20 (listing exemplary devices for output device 106).

78. The '903 patent explains that information apparatuses can have limited capabilities that prevent them from installing and/or running device drivers to process data for output:

Another challenge for mobile users is that ***many mobile information apparatuses have limited memory space, processing capacity and power.*** These limitations are more apparent for small and low-cost mobile devices including, for example, PDAs, mobile phones, screen phones, pagers, e-books, Internet Pads, Internet appliances etc. Limited memory space poses difficulties in installing and running large or complete printer or

device drivers, not to mention multiple drivers for a variety of printers and output devices. ***Slow processing speed and limited power supply create difficulties driving an output device.*** For example, processing or converting a digital document into output data by a small mobile information apparatus may be so slow that it is not suitable for productive output. ***Heavy processing may also drain or consume power or battery resources. Therefore, a method is needed so that a small mobile device, with limited processing capabilities, can still reasonably output digital content to various output devices.***

Id. at 4:11-28.

79. The '903 patent also suggests the alleged invention enables users to output content via output devices that they ***may not be able to view on their information apparatuses.***

Finally, ***some small mobile devices with limited display screens***, such as mobile phones, may in some cases be limited to ***display only a few lines of text. Browsing the Internet with such devices can be a disappointing experience*** when viewing, for example, complex web pages containing rich formats, graphics, and images. Furthermore, some small mobile devices ***may not have appropriate applications to display complex documents or languages*** such as PDF-format files, word processing documents and PowerPoint® presentation documents etc. Typically, if an application is available, displaying complex original documents on small mobile devices may require downsizing the document or page into, for example, a few lines of text. As an example, WAP protocol, I-Mode, and web clipping among others may downsize, reduce or truncate information on the original web page for display on mobile devices. Therefore, ***it is desirable to allow mobile users to output from their***

***small information apparatuses to an output device the
full richness of the original document content.***

Id. at 4:29-47; *see also id.* at 10:40-56 (describing how applications or content services may limit the types of content that can be viewed on information apparatus 100). The patent suggests that the pervasive output client capability (which involves outputting content to an output device) addresses these difficulties. *Id.* at 10:57-60 (“To address the difficulties described above, information apparatus 100 includes a pervasive output client application 102 that provides pervasive output capability of the present invention.”).

80. In my opinion, a POSA would view these discussions of information apparatuses as teaching away from using an information apparatus as an output device. In particular, it would have been clear to a POSA that devices limited memory, processing, power, and display capabilities would be unsuitable for use as output devices.

C. Subject Matter Ineligibility

81. I understand Roku asserts that the claims of the '903 patent are directed to the abstract idea of obtaining content for output. I am not a lawyer and I have not been asked to offer an opinion regarding the abstractness of the patent claims. Nonetheless, Roku's position that the claims of the '903 patent focus on the idea of obtaining content for output is consistent with my analysis of the claims.

1. Claim 1

82. Roku has asked me to analyze claim 1 of the '903 patent to provide an opinion as to whether the claim elements recite technology that was well-understood, routine and conventional. In my opinion, as explained below, claim 1 requires only technology that would have been well-understood, routine, and conventional to a POSA even as of November 1, 2000 which is the earliest claimed filing date of the '903 patent.

a. Preamble

83. The preamble of claim 1 reads as follows:

A method for outputting, at an output system, digital data content received from one or more servers over the Internet, the digital data content includes audio content or video content, the one or more servers operating, at least partly, over the Internet, the output system includes one or more devices or one or more computing devices that communicate, at least in part, with the one or more servers, by transmitting one or more objects from the output system to the one or more servers, the one or more objects being data or software entities containing information, the one or more objects being further configured to be suitable for transmitting the information from the output system to the one or more servers, the output system including:

The preamble text generally previews subject matter that is recited in more detail in the body of the claim. Rather than discuss this subject matter twice, I discuss it in the sections below.

b. Claim 1 Recites Conventional Computing Devices

84. Each of the computing devices recited in claim 1 was routine and conventional by late 2000.

(i) “Output System”

85. The '903 patent describes two basic output device configurations: (a) output devices with *internal* output controllers (*e.g.*, output device 106 in Fig. 1), and (b) output devices with *external* output controllers (*e.g.*, output device 106C and output controller 104C in Fig. 3C). '903 patent at 14:8-12 (“Output controller 104 may also be integrated, installed, or connected externally to one or more output devices 106.”); 20:54-21:44 (discussing Figs. 3A-3C which depict different output controller configurations).

86. Claim 1 appears to use the term “output system” to cover both output device configurations.⁴ Dependent claim 2 recites that the “output system” can be “a television,” for example, whereas dependent claim 3 recites that the “output system” can be a “controller box wire connected to [a] television.” Ex. 1003.

⁴ Claim 1 also recites that the “output system” includes “a wired connection to at least an output device” This requirement makes sense if the “output system” is an external output controller as recited in claim 3. The patent describes external controllers being wire connected to output devices. *See, e.g.*, '903 patent at 21:21-22 (describing configuration in Fig. 3C). The idea of an output device with an *integrated* output controller being wire connected to an output device does not seem consistent with the description in the '903 patent. There is no discussion of any output device that is wire connected to another output device.

87. Neither of the output device configurations described in the '903 patent requires any new or unconventional technology, and the patent does not suggest these configurations are themselves inventive. First, as to the “output device” aspect, the '903 patent does not describe new or improved output devices (*e.g.*, printers, monitors, televisions). The patent primarily uses printers as examples of output devices, but never states or implies that the invention is (or requires) a new or improved printer. The patent suggests, for example, that “an output controller 104 ... may be installed internally or connected externally to a *legacy printer* to provide it with wireless communication capability that was previously lacking.” *Id.* at 14:30-34. The patent also describes “conventional” printers in connection with Figures 10A and 10B. *Id.* at 5:37-41; *see also id.* at 13:33-34 (“A typical example of printer 1000B is a lower-cost inkjet printer.”). Other types of output devices such as televisions and “audio output devices” are mentioned but never meaningfully discussed. The term “television” appears three times in the written description but only in statements listing examples of output devices. '903 patent at 2:4-7; 12:1-3; 23:41-45. The '903 patent describes no new or improved television.

88. Second, the “output controller” aspect of an output system requires no new and unconventional technology either. The “output controller” described in the '903 patent is not a specific device or program. It is a black box that can be

implemented with conventional technologies in order to perform certain functions. The patent lists eleven “functionalities and components” of output controller 104, each of which is merely a function the output controller 104 may perform. *Id.* at 15:34-16:25.

89. The patent states that the output controller can be implemented with hardware, with software, or both. *Id.* at 15:5-7 (“Regardless of its implementation, the output controller 104 will usually include, hardware, software, or both.”). Rather than attempt to describe any hardware implementation, the patent merely lists well-known types of hardware components that could possibly be used in an output controller. *Id.* at 15:7-12 (“For example, an output controller 104 may include components using one or more or combinations of an application-specific integrated circuit (ASIC), a digital signal processor (DSP), a field programmable gate array (FPGA), firmware, system on a chip, and various communication chip sets.”). Similarly, the patent vaguely suggests that an output controller may contain “software components or embedded application ***software to implement its feature sets, and functionalities.***” *Id.* at 15:12-16. The implication is that a POSA would know how to design hardware and/or write software to perform the functions of the output controller without needing detailed guidance from the patent specification.

90. As I noted above, output controllers can be installed in an output

device or external. The patent describes no particular implementation of an internal output controller. It only goes as far as suggesting that an internal output controller “may be implemented as a *circuit board or card* that is installed inside an output device 106 and *may include software, hardware, or both.*” *Id.* at 14:39-42; Fig. 3A. But by 2000 it was well-known that some print servers were embodied as cards that could be installed in printers to share them on networks. *See, e.g.*, Appx. I (Neibauer, *This Wired Home*, 2000) at 245 (“For some HP LaserJet printers, you can purchase an internal print server that fits inside the printer, much the way some NICs fit inside a computer.”); Appx. H (Poole, *The Little Network Book for Windows and Macintosh*, 1999) at 130 (“Installing an internal print server usually means inserting an adapter card into the printer’s expansion slot.”). The ’903 patent suggests that a print server can be used to implement the output controller 104. ’903 patent at 14:59-61 (“Other possible implementations of output controller 104 may include, for example, a ... print server.”).

91. The ’903 patent also states that output controller 104 can be implemented with a “conventional personal computer (PC)”:


Other possible implementations of output controller 104 may include, for example, *a conventional personal computer (PC)*, a workstation, and an output server or print server. In these cases, the functionalities of output controller 104 may be implemented using application software installed in a computer (e.g., PC, server, or

workstation), with the computer connected with a wired or wireless connection to an output device 106. ***Using a PC, server, workstation, or other computer to implement the feature sets of output controller 104*** with application software ***is just another possible embodiment of the output controller 104 and in no way departs from the spirit, scope and process of the present invention.***

Id. at 14:59-15:4.

92. In view of the above, an “output system” in the context of the ’903 patent could simply be a conventional PC that serves as an output controller and outputs content over a wired connection to an output device such as a printer, monitor, or speakers.

93. Claim 1 also requires the claimed “output system” to include “one or more processors; one or more wireless communication units that include one or more chips or chipsets; an interface for interacting with a user of the output system; and a wired connection to a least an output device for outputting data content[.]” Each of these components was routine and conventional by late 2000. PCs included processors. *See, e.g.,* Appx. K (Gookin, *PCs for Dummies*, 1998) at 171-176. Conventional PCs also included “interfaces” including graphical user interfaces (GUIs) and input device interfaces (mouse, keyboard, touchpad). *Id.* at 11-12 (discussing mouse and keyboard); 77-90 (describing Windows which is GUI-based).



Id. at 11.


94. Wireless communication units were also routine and conventional by late 2000. For example, Apple launched its “AirPort”-branded line of IEEE 802.11-based wireless technologies in 1999. *See Appx. L (Apple Introduces AirPort Wireless Networking, July 21, 1999).* In its April 18, 2000 issue, PC Magazine compared six different Wireless LAN technologies: Apple AirPort, Cisco Aironet Wireless 340 Series, Compaq SL100, WL400, Lucent Orinoco Wireless Network, RadioLAN Wireless Mobilink, and 3Com AirConnect. *Appx. M (Derfler, Wireless LANs, PC Magazine, April 18, 2000).* Claim 1 requires the wireless communication unit to have “one or more chips or chipsets.” There was nothing unconventional about a wireless communication unit using chips or chipsets. The PC Magazine feature lists the chipsets for each of the wireless LAN technologies reviewed.

SUMMARY OF FEATURES						
Wireless LANs				download at www.pcmag.com		
■ YES □ NO	Apple AirPort	Cisco Aironet Wireless 340 Series	Compaq WL100, WL400	Lucent Orinoco Wireless Network	RadioLAN Wireless Mobilink	3Com AirConnect
List price per LAN adapter:						
PC Card	\$99	\$249	\$199	64-bit, \$179; 128-bit, \$199	\$299	\$219
PCI card	N/A	\$359	\$199	\$69	\$399	\$329
ISA card	N/A	\$359	N/A	\$69	\$349	N/A
List price per access point	\$299	\$1,299	\$899	\$995	\$999	\$1,195
HARDWARE						
Technology	802.11b	802.11b	802.11b	802.11b	802.3	802.11b
Chip set	Lucent	Intersil Prism II	Intersil Prism II	Lucent	PCRPIC Rev 0	Intersil Prism II

Id. at 232 (highlighting added); *see also*, e.g., Appx. N (Intersil Application Note describing Prism II WLAN chipset, February 1999). The '903 patent itself does not describe any particular chip or chipset for a wireless communication unit. The RF communication unit discussed in connection with Figure 2A is generic and described in functional terms. '903 patent at 19:32-20:19.

95. Wired connections to output devices were also routine and conventional. The specification mentions conventional serial, parallel, USB, and firewire cable connections. '903 patent at 9:67-10:2; *see*, e.g., Appx. K (Gookin, *PCs for Dummies*, 1998) at 15-16 (discussing serial, parallel (*i.e.*, “printer”), and USB connections). The specification does not mention any particular technologies for outputting to audio or video to output devices over a wired connection. Nevertheless, conventional PCs included ports (VGA out, audio out) for sending

audio and video signals to speakers and monitors.



Id. at 15-16 (highlighting added).

96. By late 2000, some computers had outputs for televisions. For example, Apple added S-Video output ports to its PowerBook G3 laptops in 1998. Appx. O (*Apple Introduces PowerBook G3*, May 6, 1998); see also Appx. P (*Macintosh PowerBook G3 Series* brochure, 1998) at 1 (“They also include built-in video output, and several configurations include S-video output for connecting to a TV.”). In September 2000, Apple released a new iBook that included an “AV port” capable of outputting a composite video signal to a TV. Appx. Q (*Apple Unveils New iBook Line*, Sept. 13, 2000) (“The new iBook models include ... an

AV port providing audio and composite video output;”); Appx. R (iBook web page archived Oct. 19, 2000) (“Well, now you can display your movies on the big screen quickly and easily, thanks to the AV port on your iBook. Simply connect your TV to the AV port using the AV cable with color-coded jacks that comes with the new iBook, and — boom — you’re on TV.”).

(ii) “One or More Servers”

97. The claimed “one or more servers” are routine and conventional too. The ’903 patent describes two servers: application server 110 and content server 114. The patent notes that these servers can be combined into a single server. ’903 patent at 19:21-27.

98. Application server 110 is described as a generic server that “may include ... mechanisms for servicing requests from a plurality of client computers.” *Id.* at 17:41-46 (“An application server 110 may include computing capability, data storage capability, and *mechanisms for servicing requests from a plurality of client computers* (referred to as clients), including the information apparatus 100, needing computational or data storage resources.”). The application server 110 may include a “processing unit, memory unit, storage unit, input/output control unit, ... a communication unit[,]” a user interface, and an operating system. *Id.* at 17:46-62. All of this is consistent with a routine and conventional application server (*e.g.*, a Windows NT server) in 2000.

99. Although the '903 patent describes application server 110 as including server application 112 that processes content into device dependent output data for a given output device (*see id.* at 18:4-9, 30:44-32:29), claim 1 does not require the “one or more servers” to perform any such processing.

100. There is nothing unconventional about content server 114 either. The patent specification states that it “may be *a web server* that hosts a plurality of web pages” that may be viewed “using an Internet browsing application such as Internet Explorer, Netscape Navigator, a WAP browser, etc.” *Id.* at 19:5-11. Alternatively, “content server 114 may be *a file server* that allows multiple clients to store and share digital files with appropriate security or authentication procedures.” *Id.* at 19:10-14 (emphasis added).

101. Web servers and file servers were routine and conventional. Web servers have been used to allow clients to retrieve content since the early 1990s. *See, e.g.,* Appx. U (Graham, *The HTML Sourcebook*, 1995) at 186-194 (describing how a client uses HTTP to retrieve a file from a web server). Windows NT Server 4 included a web server as part of its Internet Information Server (IIS). *See* Appx. G (Russel, *Running Microsoft Windows NT Server 4.0*, 1997) at 402-410. File servers are older than web servers. As an example, Sun developed the Network File System (NFS) in the 1980s which allowed client computers to use remote file

systems over a network.⁵ Windows NT Server also included file server capabilities. *See* Appx. G (Russel, *Running Microsoft Windows NT Server 4.0*, 1997) at 296 (“Microsoft Windows NT Server version 4 is a flexible network operating system that can be used in different ways [including a]s a simple data file server, with client storing their data in a central location.”). Claim 1 requires the server(s) to be accessed over the Internet. File Transfer Protocol (FTP) has long been used to allow clients to download files from servers over the Internet.

FTP (File Transfer Protocol) is a fast, efficient, and reliable way to transfer information. It was one of the first Internet services developed to allow users to transfer files from one place to another. This service is designed to let you connect your **local** machine to a **remote** computer on the Internet, browse through the files and programs that are available on the computer, and then retrieve those files to your computer.

Appx. S (Pike, *Using FTP*, 1995) at 15; *see also, e.g.*, Appx. T (Gralla, *How the Internet Works*, 1998) at 127 (“ONE of the most popular uses of the Internet is to download files—that is, transfer files from a computer on the Internet to your computer. These files can be of many types: programs that you can run on your own computer; graphics you can view; sounds and music you can listen to; or text files that you can read. Many tens of thousands of files are downloaded every day on the Internet. Most of those files are downloaded using the Internet’s File

⁵ *See, e.g.*, <https://datatracker.ietf.org/doc/html/rfc1094> (RFC 1094: *NFS Network File System Protocol Specification*, March 1989).

Transfer Protocol, commonly referred to as FTP.’). Windows NT’s Internet Information Server included FTP server capabilities. *See* Appx. G (Russel, *Running Microsoft Windows NT Server 4.0*, 1997) at 408-410.

(iii) “Client Device”

102. The claimed “client devices” are also conventional and generic. The term “client device” appears in just one paragraph in the ’903 patent and is used to describe information apparatus 100 when it interacts with the application server. ’903 patent at 31:53-32:7. The client devices / information apparatuses can be “without limitation, desktop computers, laptop computers, networked computers, palmtop computers (hand-held computers), personal digital assistants (PDAs), Internet enabled mobile phones, smart phones, pagers, digital capturing devices (e.g., digital cameras and video cameras), Internet appliances, e-books, information pads, and digital or web pads.” *Id.* at 1:61-2:7. The patent does not contain any meaningful description of any of these devices or how the differences between them impacts the invention in any way.

103. The ’903 patent’s discussion of information apparatuses (e.g., information apparatus 100) focuses not so much on the devices themselves as on the functionality provided by “pervasive output client application 102” on the devices. Even as to that application, the patent states that it “may be variously implemented” and describes it in functional terms. *Id.* at 10:57-11:61 (listing 11

“functionalities”); 25:23-29:44 (describing “exemplary client application process 601” as shown in Figure 6). Much of the functionality described for client application 102 (*e.g.*, obtaining the various objects, sending the objects to the server, receiving processed output data) is not even required of the “client device” in claim 1.

c. Claim 1 Recites Conventional Wireless Communications

104. The '903 patent does not describe any new or improved wireless communication technologies. The specification states that Bluetooth, HomeRF, and IEEE 802.11 “all provide solutions for wireless local area networks (LANs).” 903 patent at 19:56-63. The patent also states that “[i]t is anticipated that new local area wireless technologies may emerge or that the existing ones may converge. Nevertheless, all theses [sic] existing and future wireless technologies may be implemented in the present invention without limitation.” *Id.* at 19:64-20:3; *see also id.* at 9:59-61 (“As an example, the wireless interface may be a short-range radio interface such as those implemented according to the Bluetooth or IEEE 802.11 standard.”). The patent thus embraces the use of standard wireless technologies such as IEEE 802.11 and Bluetooth.

d. Claim 1 Recites Conventional Steps

105. Claim 1 recites “[a] method for outputting ... digital data content” that includes 12 steps. There are ten enumerated steps and then two additional steps

recited in elements 1[q] and 1[r]. Ex. 1003. As discussed below, each of the steps requires only computer functionality that was routine and conventional.

(i) Step 1: Obtaining Authentication Information

106. Element 1[f] requires obtaining authentication information for accessing a service that provides content:

(1) obtaining, by the output system, authentication information for accessing a service provided by the one or more servers, the service includes providing, by the one or more servers and to the output system, one or more digital content that is available at the one or more servers, for outputting at the output system, the authentication information being related to the output system or the user of the output system;

107. The '903 patent gives examples of “authentication information such as a password, user name, id number, biometric information, digital certificate or security key.” ’903 patent at 27:7-12; *see also id.* at 27:39-42 (describing similar authentication information in context of authentication by output device). This authentication information may be manually input by users according to the patent. *Id.* at 27:14-16 (“[S]erver application 112 may ***prompt the user to enter authentication information*** through a GUI in information apparatus 100 when necessary.”); 29:20-22 (“user may be prompted to log in or sign up when he or she requests service”); 31:48-52 (“the user may input authentication information as needed”).

108. By November 2000, there was nothing unconventional about

obtaining authentication information (*e.g.*, username and password) from a user.

FTP servers have long been able to restrict access to authenticated users. For example, an FTP specification published in 1985 describes a user inputting a username and password for authentication by the server:

7. TYPICAL FTP SCENARIO

User at host U wanting to transfer files to/from host S:

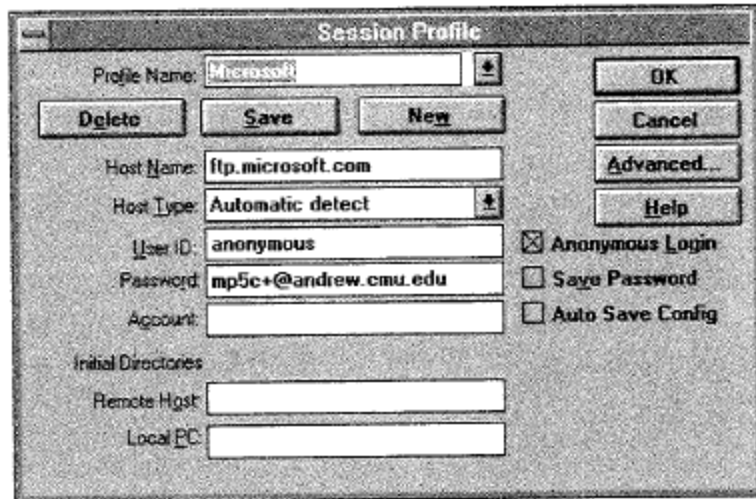
In general, the user will communicate to the server via a mediating user-FTP process. The following may be a typical scenario. The user-FTP prompts are shown in parentheses, '---->' represents commands from host U to host S, and '<----' represents replies from host S to host U.

LOCAL COMMANDS BY USER	ACTION INVOLVED
ftp (host) multics<CR>	Connect to host S, port L, establishing control connections. <---- 220 Service ready <CRLF>.
username Doe <CR>	USER Doe<CRLF>----> <---- 331 User name ok, need password<CRLF>.
password mumble <CR>	PASS mumble<CRLF>----> <---- 230 User logged in<CRLF>.

Appx. V (RFC 959) at 59; *see also id.* at 25-26 (discussing username and password commands). Popular GUI-based FTP clients such as WS_FTP allowed users to input their usernames and passwords to connect to restricted servers:

Fig. 5.2

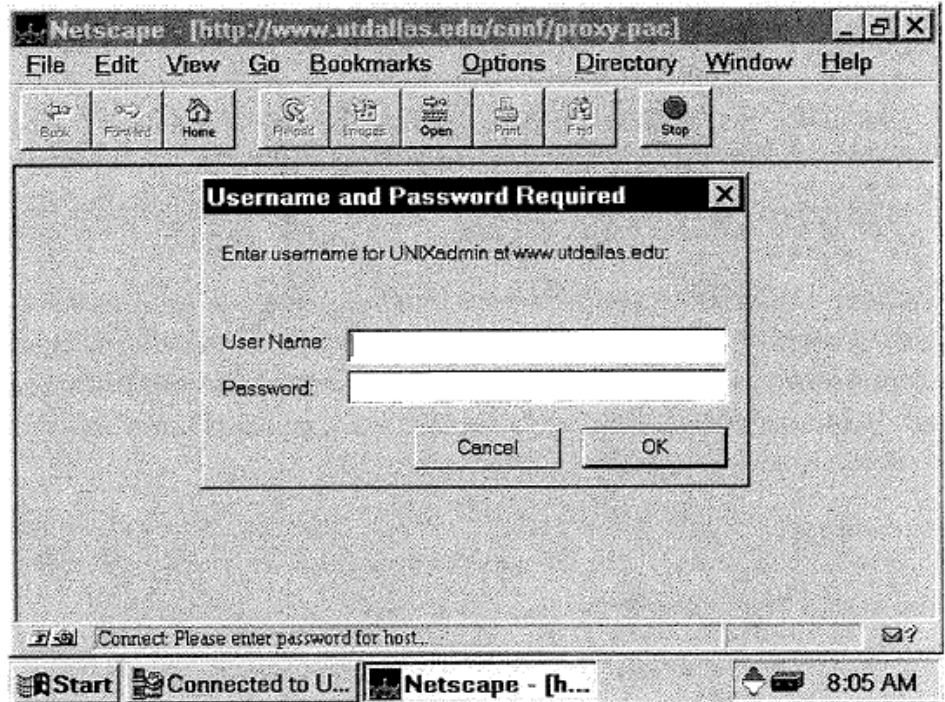
The WS_FTP Session Profile dialog box lets you connect to any FTP server on the Internet.



Appx. S (Pike, *Using FTP*, 1995) at 79; *see also id.* at 46-47 (showing and describing dialog box with user name and password fields for connecting to an FTP server); *see also, e.g.,* Appx. T (Gralla, *How the Internet Works*, 1998) at 127 (“To log on to an FTP site and download files, an account number (or username) and a password must be typed in before the daemon will allow you to enter.”).

109. Web browsers also obtained authentication information from users to access restricted web sites. The HTTP 1.0 specification, published in 1996, describes an “Access Authentication” mechanism “which may be used by a server to challenge a client request and by a client to provide authentication information.” Appx. W (RFC 1945) at 47. To enable users to access restricted sites, browsers obtained credentials from users and sent them to the server. *See id.* at 48-49. The figure below shows a basic authentication login screen in the Netscape browser:

FIGURE 46.4.
*A basic authentication
login screen.*



Appx. X (Ellsworth, *Internet 1997 Unleashed*) at 827; see also, e.g., Appx. Y (Rubin, *Web Security Sourcebook*, 1997) at 73 (“Many Web sites require their users to register a name and a password. When users connect to these sites, their browser pops up an authentication window that asks for these two items.”). A more secure authentication scheme, known as digest authentication, was standardized for use with HTTP. Digest access authentication avoids sending passwords in clear text by sending a checksum (or “digest”) generated from certain data including the password. Digest access authentication is described in RFC 2069⁶ published in January 1997 and RFC 2617⁷ published in June 1999. Both

⁶ Available at: <https://datatracker.ietf.org/doc/html/rfc2069>.

⁷ Available at <https://datatracker.ietf.org/doc/html/rfc2617>.

authentication schemes were well-understood. *See, e.g.,* Appx. Y (Rubin, *Web Security Sourcebook*, 1997) at 144-150 (discussing both basic and digest authentication).

110. Login forms embedded within web pages were another routine and conventional way of obtaining authentication information from a user. *See* Appx. U (Graham, *The HTML Sourcebook*, 1995) at 183 (“Suppose a user retrieves a fill-in HTML FORM from a server and enters his or her username and password information to access a restricted server-side resource. When the user submits the FORM data to the server, this username/password information is sent to the server gateway program as part of the data contents”).

(ii) Step 2: Connecting to a Wireless LAN

111. Element 1[g] requires “wirelessly coupling” the output system to a wireless LAN in a way that is compatible at least in part with IEEE 802.11.

(2) wirelessly coupling, by the output system and using at least one chip or chipset of the one or more chips or chipsets of the output system, the output system to a wireless local area network, wherein the at least one chip or chipset of the output system is compatible, at least in part, with at least part of a protocol within IEEE 802.11 wireless standards for coupling the output system to the wireless local area network;

112. The '903 patent describes no new or different way of coupling a device to a wireless LAN. It merely states that “implementations based on IEEE 802.11 standard ... provide solutions for wireless local area networks (LANs).”

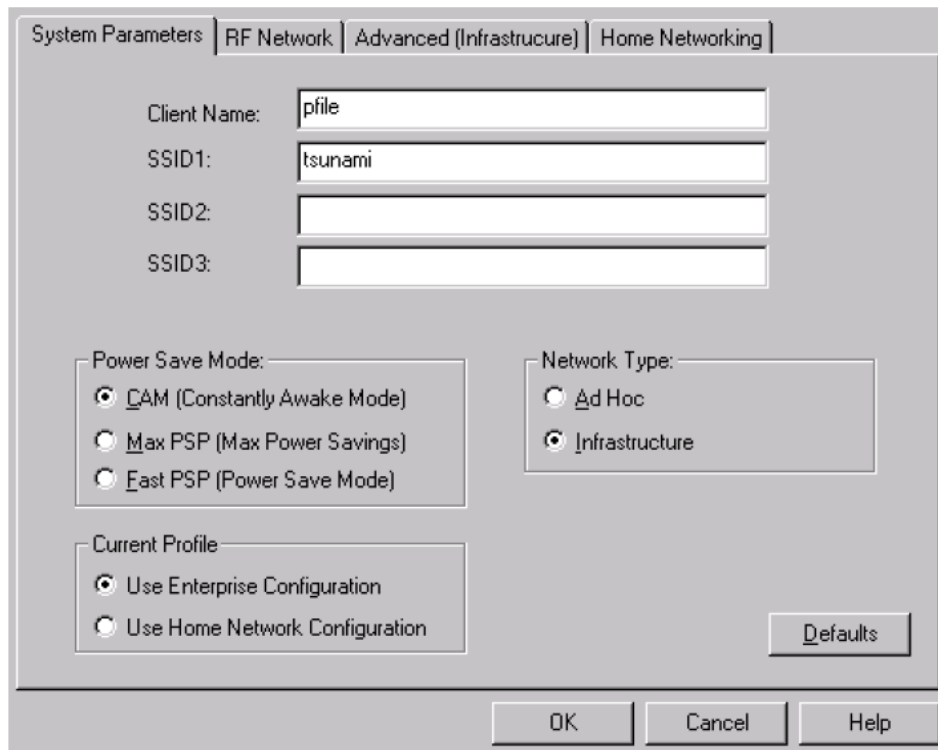
'903 patent at 19:58-63.

113. By November 2000, there was nothing unconventional about connecting a device to a wireless LAN. Wireless modules often came with associated software for connecting to wireless networks. For example, Apple provided software in connection with its AirPort wireless cards including an “AirPort Control Strip” utility that could be used to connect to an available network. The figure below shows use of the control strip to connect to a wireless network.



Appx. Z (Pogue, *The iBook for Dummies*, 2000) at 288. The figure below shows a configuration utility for the Cisco's 802.11 Aironet 340 client adapter which allows the user to enter one or more SSIDs (Service Set Identifiers) to connect to wireless networks:

Figure 2 Each Aironet 340 series client adapter comes complete with a powerful, yet intuitive, Windows-based configuration, management, and diagnostics utility.



Appx. AA (*Cisco Aironet 340 Series Client Adapters and Access Points*, 2000) at 3; *id.* (“To further facilitate installation, Cisco provides a suite of integrated utilities for Windows-based configuration, management, and diagnostics.”).

(iii) Step 3: Connecting to a Server on the Internet

114. Element 1[h] requires wirelessly connecting to the server(s) providing the content service:

(3) wirelessly connecting the output system, by the output system, using the at least one chip or chipset of the output system, and over the wireless local area network wirelessly coupled in (2), to the one or more servers over the Internet;

115. The '903 patent states that “client application 102 may communicate with the server application 112 using one or more or a combination of *standard*

network protocols such as WAP, ... **TCP/IP**, ... NetBEUI, Apple Talk, among others.” 28:56-61. The patent further states that “[t]he communication link between information apparatus 100 and application server 112 may be implemented with one or a combination of **standard network connections and communication links**” *Id.* at 28:62-29:2.

116. As I discussed above, the ’903 patent suggests that content server 114 can be implemented as a web server or file server that can be accessed with a web browser or other “appropriate application.” *Id.* at 19:1-20. No specifics are provided as to how information apparatus 100 or any other device connects to a content server.

117. There was nothing unconventional about connecting to a remote server in 2000. Conventional web browsing involves clients (*e.g.*, web browsers) making TCP connections to web servers and using HTTP to request resources from the server. Appx. BB (Stevens, *TCP/IP Illustrated, Vol. 3*, 1996) at 161-162. The figure below illustrates the TCP connections between a client and three web servers:




Figure 13.2 Organization of a Web client-server.

Id. at 162. *See also* Appx. U (Graham, *The HTML Sourcebook*, 1995) at 183 (discussing the opening of a connection prior to making an HTTP request).

118. FTP has also long been used to connect to servers on the Internet. *See, e.g.,* Appx. S (Pike, *Using FTP*, 1995) at 15 (“FTP (File Transfer Protocol) is a fast, efficient, and reliable way to transfer information. It was ***one of the first Internet services developed to allow users to transfer files from one place to another.*** This service is designed to let you connect your local machine to a ***remote computer on the Internet***, browse through the files and programs that are available on the computer, and then retrieve those files to your computer.”).

119. Similar to HTTP, FTP runs on top of TCP and uses TCP to establish connections between FTP clients and servers. *See, e.g.,* Appx. V (RFC 959, October 1985) at 59 (“8. CONNECTION ESTABLISHMENT ... The FTP control connection is established via TCP between the user process port U and the server process port L.”); *see also* Appx. T (Gralla, *How the Internet Works*, 1998) at 128-

129 (showing and describing the “command link” and “data link” connections used by FTP).

(iv) Step 4: Sending Authentication Information

120. Element 1[i] requires wirelessly sending a job object with the authentication information obtained in element 1[f]:

(4) wirelessly sending, by the output system, using the at least one chip or chipset of the output system, and over the wireless local area network wirelessly coupled in (2), a job object, which includes the authentication information obtained by the output system in (1), from the output system to at least one server of the one or more servers over the Internet for accessing, by the output system, the service provided by the one or more servers, the job object being an object from among the one or more objects, and the job object being an object that is related to an output job;

121. The '903 patent specification makes clear that the “job object” merely refers to certain data sent to the server. The patent uses the term “object” to refer to abstract entities that may contain data and/or software. '903 patent at 5:55-57 (“An object may refer to a software and data entity, which may reside in different hardware environments or platforms or applications.”). The patent does not suggest that an object needs to be implemented in any particular way:

It is important to note that the term object is not limited to software or data as its media. *Any entity containing information, descriptions, attributes, data*, instructions etc. *in any computer-readable form or medium* such as hardware, software, files based on or including voice,

text, graphics, image, or video information, etc., are all
valid forms of object definition.

Id. at 5:64-6:3. The patent also suggests that the “objects” can consist of a single attribute. *Id.* at 6:43-46 (“An output device object may contain ***one or more*** attributes that may identify and describe, for example, the capabilities and functionalities of a particular output device such as a printer.”); 7:52-54 (“Examples of attributes and information contained in a job object may include ***one or more*** of the following, among others.”); 8:10-12 (“A document object may contain ***one or more*** of the following attributes, fields, or descriptions.”).

122. The patent states that “[a] job object may contain attributes and information that describe an output job.” *Id.* at 7:43-44. The patent identifies nine categories of information that could be included in a job object, none of which is strictly required. *Id.* at 7:52-8:6. The second category is “[i]nformation on security, authentication, payment, subscription, identification among others.” *Id.* at 7:61-62. Claim 1 does not require the claimed job object to include anything other than the authentication information.

123. The ’903 patent describes no new or improved way of sending a job object to a server. The patent generally states that “transferring an object over the network ***may involve protocols such as file transfer protocol (FTP) or hypertext transfer protocol (http)***, among others.” ’903 patent at 6:31-33. The patent also describes how the client application 102 on information apparatus 100 sends the

objects to the server in a “composite message.” *Id.* at 28:48-54. The patent explains that this “composite message may be *any type of data transferred across network 108 that may include one or more transmissions.*” *Id.* at 28:49-51. The patent also suggests that the composite message can be sent “using one or more or a combination of *standard networking protocols*” and using “*standard network connections and communication links.*” *Id.* at 28:55-29:2.

124. By late 2000, there was nothing unconventional about sending authentication information to a server. As I discussed above, FTP users often provided usernames and passwords when connecting to FTP servers. Even FTP servers that allow “anonymous” access often required users to provide a username (*i.e.*, “anonymous”) and a password (*e.g.*, the user’s email address). *See* Appx. S (Pike, *Using FTP*, 1995) at 81. Telnet is another old and well-known application protocol that allows computers to access a remote host over a network.⁸ It was very common to use telnet to connect to a server or other remote workstation and then login to the remote computer with a username and password in order to use its resources.

125. Sending authentication information in order to access content from a web server was also routine and conventional. I discussed the HTTP basic and

⁸ *See, e.g.*, <https://datatracker.ietf.org/doc/html/rfc854> (RFC 854: *Telnet Protocol Specification*, May 1983).

digest access authentication schemes above. With the basic scheme, the username and password (authentication information) were sent to the server in an encoded string. Appx. W (RFC 1945, *Hypertext Transfer Protocol -- HTTP/1.0*, May 1996) at 48-49 (“To receive authorization, the client sends the user-ID and password, separated by a single colon (‘:’) character, within a base65 [5] encoded string in the credentials.”). With the digest scheme, a more secure “digest” is sent and that digest is generated using the user’s password. Appx. CC (RFC 2069, *An Extension to HTTP: Digest Access Authentication*, January 1997) at 6-7. The digest is “authentication information” because the server compares it to a digest it computes itself to determine whether to authorize the user. *Id.* at 10.

126. As I also mentioned above, it was also common for websites to include login forms in their web pages. When the user submitted the form, their credentials were sent to the server so they could be compared to authentication information stored for registered users. *See, e.g.,* Appx. U (Graham, *The HTML Sourcebook*, 1995) at 183-184 (“Suppose a user retrieves a fill-in HTML FORM from a server and enters his or her username and password information to access a restricted server-side resource. When the user submits the FORM data to the server, this username/password information is sent to the server gateway program as part of the data contents”).

(v) Step 5: Accessing a Content Service

127. Element 1[j] requires wirelessly accessing the content service based on having sent the job object / authentication information:

(5) wirelessly accessing, by the output system, using the at least one chip or chipset of the output system, and over the wireless local area network wirelessly coupled in (2), the service provided by the one or more servers over the Internet, the wirelessly accessing of the service being based, at least in part, on the output system having wirelessly sent the job object to the at least one server of the one or more servers in (4);

128. As discussed, the '903 patent describes content server 114 which may be a web server or file server. Where a web server is used, the user accesses its pages with a web browser. '903 patent at 19:5-11. Where a file server is used, the patent states that it may support “appropriate security or authentication procedures” and that “a user may need to have appropriate application on his/her information apparatus 100 to access, view, and edit these files.” *Id.* at 19:10-20. There is no more detailed technical description of how any device “accesses” any service provided by a server.

129. There was nothing unconventional about accessing a content service provided by a server over a network in 2000. As the '903 patent itself suggests, websites are content services that respond to user requests by providing web pages and other data. *Id.* at 19:5-11. As I have discussed in addressing Steps (1) and (4), the HTTP protocol provided ways for web servers to restrict access to authorized

users that have been authenticated. *See, e.g.*, Appx. W (RFC 1945, May 1996) at 47 (“HTTP provides a simple challenge-response authentication mechanisms which may be used by a server to challenge a client request and by a client to provide authentication information.”).

130. FTP sites were also examples of content services that were accessed by authenticated users over the Internet. *See, e.g.*, Appx. S (Pike, *Using FTP*, 1995) at 15 (describing FTP as a “**service** designed to let you **connect your local machine to a remote computer on the Internet**, browse through the files and programs that are available on the computer, and then retrieve those files to your computer”). As I have discussed above, FTP servers could restrict access to authenticated users.

(vi) Step 6: Receiving User Selection of Content

131. Element 1[k] requires receiving an indication of selected digital content from among the content available at the server:

(6) receiving, by the output system and via the interface of the output system, an indication of selected digital content from among the one or more digital content that are available at the one or more servers;

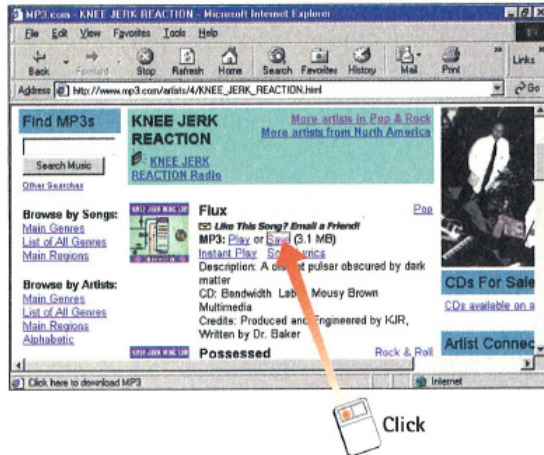
132. As I discussed above, '903 patent describes a user identifying selected content using a GUI within client application 102 or through interactions with another application “(e.g., document editing and or browsing application).” *See supra* ¶¶ 56-58.

133. By November 2000, there was nothing unconventional about receiving a user selection of digital content. For example, when users surf the web, they select hyperlinks to navigate to new pages. Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 583 (“Links among pages, shown in [Figure 35.1] as directed arrows, connect an anchor on one page of hypertext to another hypertext page or a specific location on that page.... Anchors in hypertext are displayed as hotspots in a Web browser and are often shown as highlighted or underlined (or both) text that the user can select, often using a point-and-click interface.”); 626 (“By clicking links, you can travel from page to page, from one country to another, anywhere in the world ...”). Receiving selection of a hyperlink is “receiving ... an indication of selected digital content” as claimed. There were other well-understood ways for users to identify selected content including manually typing in a URL in their web browser or selecting a previously saved bookmark.

134. In addition to web pages, hyperlinks can point to files, including audio or video files, to be downloaded. Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 583 (“Hypertext links can also link to multimedia—sound, graphics, movies, or interactive content.”); *see also id.* at 608-610 (downloading linked files in Netscape), 635-636 (downloading linked files in Internet Explorer). Popular websites like MP3.com provided links to download MP3 audio files:

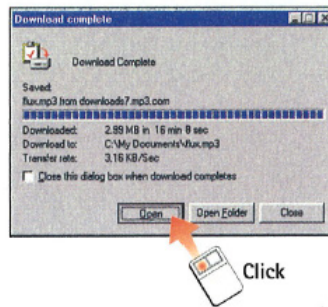
4 Download a Song

When you find a song you'd like to listen to, click the **Save** hyperlink to open a dialog box that lets you choose where to store the song on your system. A dialog box with a progress bar opens to track the download of that file to your system.



5 Open a Song File

When the MP3 file has finished downloading, click the **Open** button to play it.



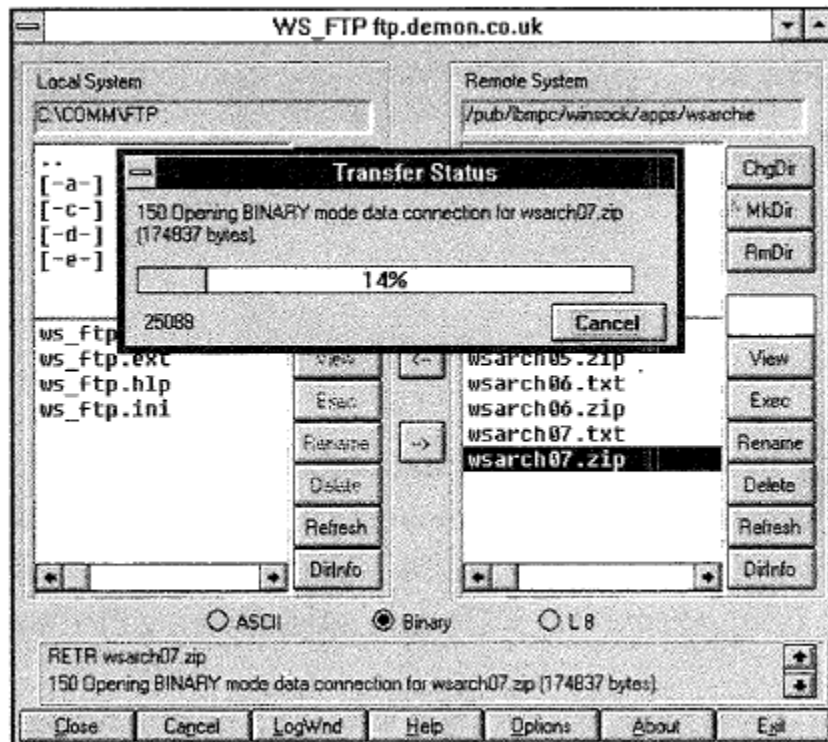
Appx. DD (Cadenhead, *How to Use the Internet*, 2000) at 147. *See, also, e.g.*, Appx. EE (Musciano, *HTML & XHTML*, 2000) at 168 (“You normally use the anchor tag (<a>) to link external multimedia elements to the current document. Just like other link elements selected by the user, the browser downloads the multimedia object and presents it to the user, possibly with the assistance of an external application or plug-in.”); *id.* at 169 (“Just like any referenced document, the server delivers the desired multimedia object to the browser when the user selects the link. If the browser finds the document is not HTML or XHTML, but some other format, it automatically invokes an appropriate rendering tool to display or otherwise convey the contents of the object to the user. ... If a browser has not been configured to handle a particular document format, the browser will

inform you and offer to simply save the document to disk.”).

135. FTP also involves receiving indications from users. In a GUI-based FTP client like WS_FTP, the user can simply double-click the file they wish to download or select the file and then select a download button:

Fig. 6.1

After double-clicking the file you want to download, WS_FTP notifies you of its progress as it transfers the file to your default directory.



Appx. S (Pike, *Using FTP*, 1995) at 93-94 (“Simply double click the file you want to retrieve. (Or click the file to select it, and then click the left arrow in the middle of the screen.)”).

136. In command-line FTP clients, the user indicates selection of a digital document by inputting its name as part of a command (e.g., a “get” command) to retrieve the file. Appx. S (Pike, *Using FTP*, 1995) at 222 (“Simply follow the get command with the *name of the file you want* (and, optionally, with the name you

want to store the file under on your local account.) For example **get pkz204g.exe**
pkz.exe moves the file pkz204g.exe from the FTP server to a file called pkz.exe on
your local account.”).

**(vii) Step 7: Identifying Selected Content to the
Server**

137. Element 1[l] requires sending a document object to the server that
identifies the content selected in step 6:

(7) wirelessly sending, by the output system and using
the at least one chip or chipset of the output system, from
the output system, over the wireless local area network
wirelessly coupled in (2), a digital document object to at
least one server of the one or more servers over the
Internet, the digital document object includes at least a
pointer or a reference to the digital content selected in
(6), the digital document object being an object from
among the one or more objects, and the document object
being an object that is related to the digital content
selected in (6);

138. The '903 patent suggests that objects can reference content using
information such as file names and URLs. '903 patent at 6:11-15 (“Examples of
reference may include universal resource identifier scheme (URI), uniform
resource locator (URL), IP address, file names, directory pointers, software object
and component pointers, and run-time address, among others.”). Because a
document object could consist of a single attribute (*id.* at 8:28-29), a document
object could be as simple as just a URL or just a file name. The patent also
describes a user inputting a file name or URL associated with a remotely stored

digital document. *Id.* at 25:51-55 (“[T]he client application 102 may provide a GUI with which a user can directly input the pointer or reference (e.g., URL, IP address, filename, path, etc.) of a digital document stored locally or in a network node.”).

139. The ’903 patent describes sending a document object to the server as part of a “composite message.” ’903 patent at 28:51-54. As I have noted, that “composite message may be any type of data transferred across network 108” and may be sent with “standard network protocols” using “standard network connections and communication links.” *Id.* at 28:48-29:2. The patent also generally states that protocols such as HTTP and FTP can be used for object transmission. *Id.* at 6:31-33.

140. By late 2000, there was nothing unconventional about sending data to a server that identifies a content item selected by a user. When a user selects a hyperlink in a web browser, the URL referenced by that hyperlink (*i.e.*, its HREF attribute) is transmitted to the server as part of an HTTP GET request. *See, e.g.*, Appx. U (Graham, *The HTML Sourcebook*, 1995) at 187 (“The request message contains ... the *method* field ... which specifies the HTTP method to be used and ***the location of the desired resource on the server***”); *see also generally id.* at 186-194 (describing as “Example 1” a simple GET method request that “could be initiated by clicking on a hypertext anchor pointing to the file.”). As I noted

above, the URL referenced by a hyperlink can point to a web page or a file of another type to be downloaded. Additionally, there were other ways to transmit information from clients to web servers. It was common to append information to a URL included in a GET request, for example. *See, e.g., id.* at 194-198 (describing HTTP GET request with search strings and form submission using GET). Additionally, HTTP supported a POST method that allowed clients to submit information in the body of an HTTP request rather than in the URL. *See, e.g., id.* at 198-200 (describing form submission using POST).

141. Similarly, when a user retrieves a file from an FTP server, the file selected by the user is identified to the server. As I discussed, in a command-line utility the user enters a command like “get” followed by the name of the file to be retrieved. Appx. S (Pike, *Using FTP*, 1995) at 222. GUI-based FTP clients generate the same kinds of commands to retrieve files from FTP servers. *See id.* at 90 (“You can see the actual FTP commands executed for any remote operation in the log area at the bottom of the WS_FTP window.”); *id.* (sample log showing “RETR gw2000-FAQ” command leads to transfer of gw2000-FAQ file).

(viii) Step 8: Receiving Content

142. Element 1[m] requires receiving output data related to the document selected in step 6:

(8) wirelessly receiving, by the output system, using the at least one chip or chipset of the output system and over

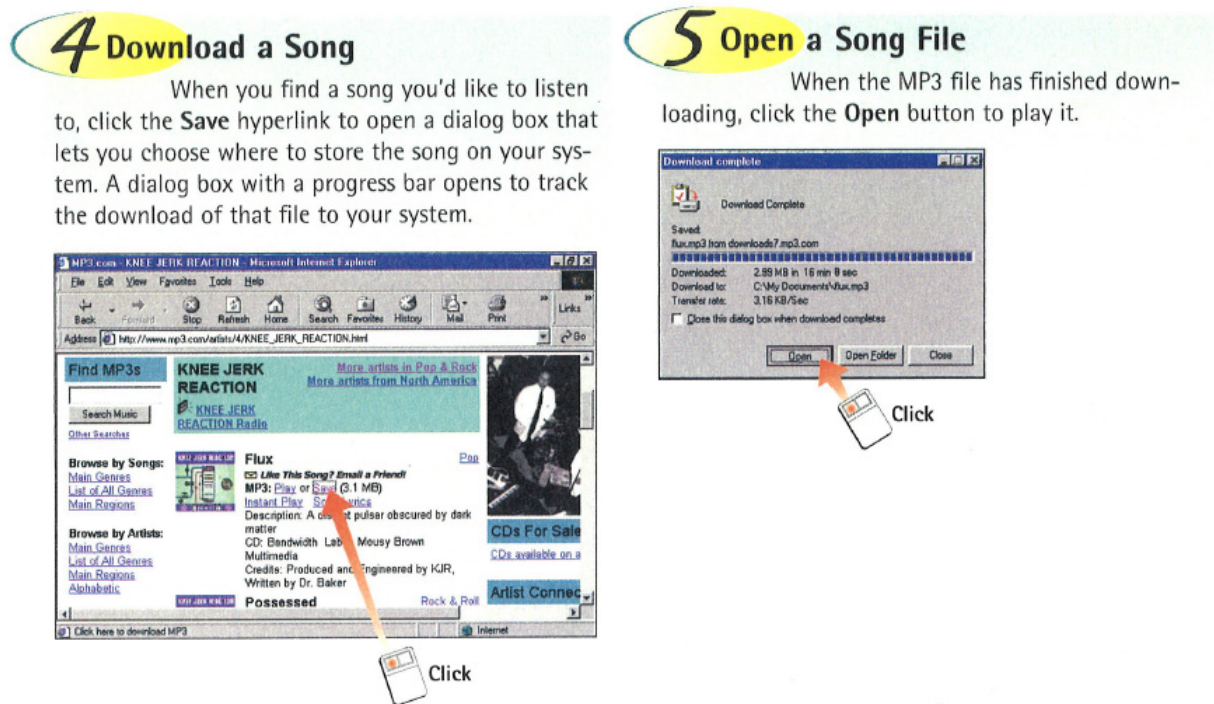
the wireless local area network wirelessly coupled in (2), output data from at least one server of the one or more servers over the Internet, the wireless receiving of the output data being subsequent to the output system having wirelessly sent the digital document object in (7) to the at least one server of the one or more servers over the Internet, and at least part of the output data includes audio digital content or video digital content, individually or in any combination, which is related to the digital content selected in (6), and the receiving of the output data being subsequent to the output system having wirelessly sent the job object in (4) to the at least one server of the one or more servers over the Internet;

143. The '903 patent states that the server “transmits output data to the information apparatus 1000 through network 108.” '903 patent at 32:30-31; *see also id.* at 23:3-5; 29:63-65. No new or improved way of receiving the output data is described.

144. By late 2000, there was nothing unconventional about receiving selected content from a server over the Internet. Conventionally, when a user selects a hyperlink to a file in a web browser, the file is downloaded (*i.e.*, received) to the user's computer. *See, e.g.*, Appx. EE (Musciano, *HTML & XHTML*, 2000) at 168 (“You normally use the anchor tag (<a>) to link external multimedia elements to the current document. Just like other link elements selected by the user, the **browser downloads the multimedia object** and presents it to the user, possibly with the assistance of an external application or plug-in.”); *id.* at 169 (“Just like any referenced document, the server delivers the desired multimedia

object to the browser when the user selects the link. If the browser finds the document is not HTML or XHTML, but some other format, it automatically invokes an appropriate rendering tool to display or otherwise convey the contents of the object to the user. ... If a browser has not been configured to handle a particular document format, the browser will inform you and offer to simply save the document to disk.”).

145. The figure below shows the downloading of an MP3 file from MP3.com after the user selects the “Save” hyperlink on the webpage.



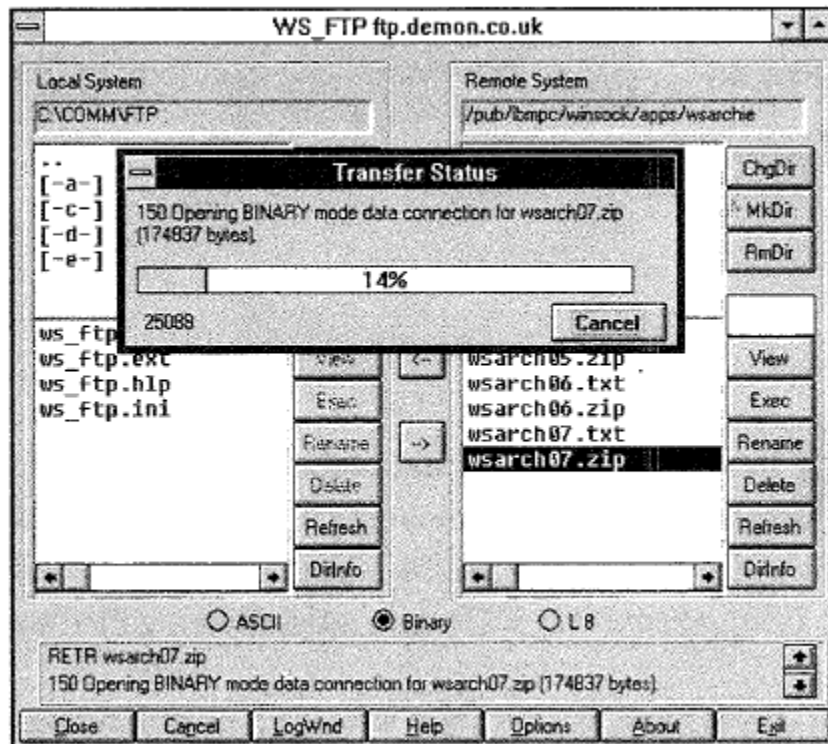
Appx. DD (Cadenhead, *How to Use the Internet*, 2000) at 147.

146. Similarly, when a user retrieves a file via FTP, they receive the file from the FTP server. The figure below shows the receiving of a file requested by

the user in the WS_FTP client.

Fig. 6.1

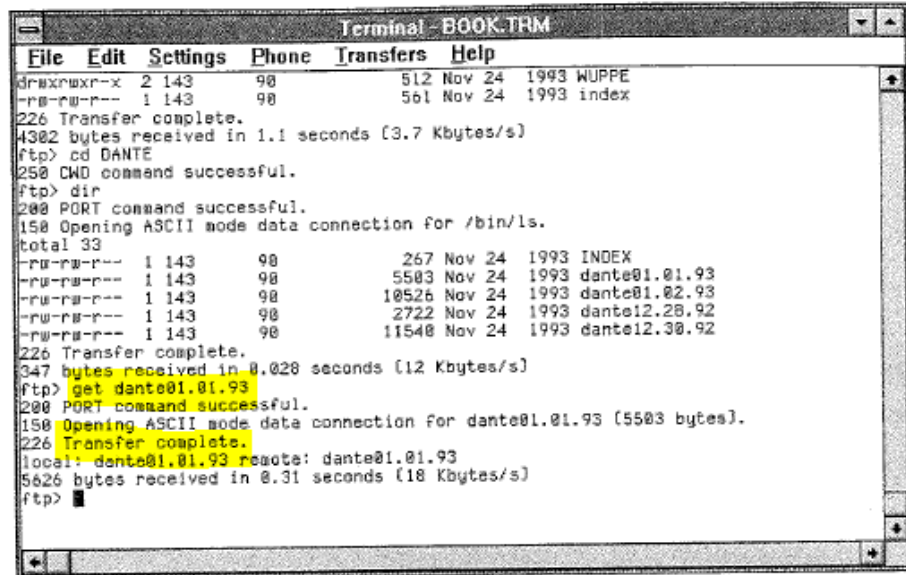
After double-clicking the file you want to download, WS_FTP notifies you of its progress as it transfers the file to your default directory.



Appx. S (Pike, *Using FTP*, 1995) at 94. The figure below similarly shows successful receipt of a file after issuing a “get” command for the file using a command-line based FTP client:

Fig. 14.5

When ftp finishes transferring a file, it tells you the name of the local and remote file, the size of the file in bytes, and how long the transfer took.



```
Terminal - BOOK.TRM
File Edit Settings Phone Transfers Help
drwxrwxr-x 2 143 98 512 Nov 24 1993 WUPPE
-rw-rw-r-- 1 143 98 561 Nov 24 1993 index
226 Transfer complete.
4302 bytes received in 1.1 seconds (3.7 Kbytes/s)
ftp> cd DANTE
250 CWD command successful.
ftp> dir
200 PORT command successful.
150 Opening ASCII mode data connection for /bin/ls.
total 33
-rw-rw-r-- 1 143 98 267 Nov 24 1993 INDEX
-rw-rw-r-- 1 143 98 5503 Nov 24 1993 dante01.01.93
-rw-rw-r-- 1 143 98 10526 Nov 24 1993 dante01.02.93
-rw-rw-r-- 1 143 98 2722 Nov 24 1993 dante12.28.92
-rw-rw-r-- 1 143 98 11540 Nov 24 1993 dante12.30.92
226 Transfer complete.
347 bytes received in 0.020 seconds (12 Kbytes/s)
ftp> get dante01.01.93
200 PORT command successful.
150 Opening ASCII mode data connection for dante01.01.93 (5503 bytes).
226 Transfer complete.
local: dante01.01.93 remote: dante01.01.93
5626 bytes received in 0.31 seconds (18 Kbytes/s)
ftp>
```

Id. at 223 (highlighting added).

(ix) Step 9: Processing Content

147. Element 1[n] requires processing part of the output data that involves an operation related to any of 16 different possible processing steps:

(9) processing, at the output system and using the one or more processors of the output system, at least part of the output data wirelessly received in (8) into audio output data or video output data for outputting or playing at least part of the digital content selected in (6) at the output system or at the at least an output device, the processing of the at least part of the output data includes one or more operations related to a decoding operation, an encoding operation, an encryption operation, a decryption operation, a compression operation, a decompression operation, a conversion operation, an image enhancement operation, an image processing operation, a color correction operation, a color management operation, an interpolation operation, a scaling operation, a smoothing operation, a segmentation operation, or a de-segmentation operation, individually or in any combination; and

148. The 16 processing operations mentioned in step (9) were well-known and conventional by the time of the alleged invention. The '903 patent itself describes many of the recited processing steps as being performed by the “conventional” printer 1000A in Fig. 10A. '903 patent at 12:58-62 (raster image processor 1002 in printer 1000A may perform “scaling, segmentation, color space transformation, image enhancement, color correction, halftoning, compression, etc.”); 13:9-15 (“Some printer manufacturers may also employ ... encoding, decoding, compression, decompression, etc. for the print data.”); 5:36-38 (identifying printer in Fig. 10A as “conventional”). The '903 patent does not describe improved processing operations (*e.g.*, improved compression / decompression) or suggest that the alleged invention changes the way that these known processing operations are performed at the output system. The processing of the output data is described only at a very high level. *Id.* at 32:49-58 (listing optional processing steps that may be performed by information apparatus 100 after receiving output data from the server); Fig. 8 (step 804); *see also id.* at 33:5-10 (“output controller 104 or combined controller ... may ... further process the data”).

149. Even though the '903 patent's claims require processing the output data into ***audio and video output data***, there is no discussion in the patent of processing audio or video data in particular or how such processing might differ

from processing document data for output by a printer.

150. There was nothing unconventional about processing output data with one or more of the 16 generic processing operations listed in the claim. Processing operations such as “decoding,” “decryption,” “decompression,” and “conversion” were well-understood and widely used by 2000.

151. Because compressing files can make them smaller and reduce the time and bandwidth required to transfer them, files available for download from FTP servers are often compressed. *See, e.g.,* Appx. S (Pike, *Using FTP*, 1995) at 60 (“FTP sites use compression and encoding for a variety of reasons. To try to solve the dilemma of file size, FTP sites use compression. Compression makes a binary file smaller so it can be stored and retrieved more efficiently. A compressed file can’t be executed—it must be decompressed first.”); Appx. T (Gralla, *How the Internet Works*, 1998) at 127 (“As a way to speed up file transfers and save space on the FTP server, files are commonly compressed After the files have been downloaded, you’ll need to run the compression software on your own computer to decompress the files so you can use them.”).. Programs like WinZip or PKUNZIP could be used to decompress .zip files, for example. Appx. S (Pike, *Using FTP*, 1995) at 61. Some FTP clients can automatically decompress and/or decode downloaded files. *See id.* at 39 (describing Fetch).

152. It was also routine and conventional for files available for download

from the web to be compressed. *See, e.g.,* Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 608 (“Many Web pages offer links to files that can’t be viewed by Netscape directly, such as Windows or Mac programs or ***compressed archives of data.***”); *id.* at 666-667 (“It has long been the case that files made available for download have been ***stored in compressed formats***, thus reducing the amount of time needed to transfer information via modem.... Needless to say, one of the most basic tools to have when browsing the Web is a program to decompress these archived files. For Windows, a hand tool to use is WinZip, based on PKZip.”); *id.* at 667 (discussing decompression utilities for Mac). Downloaded files compressed with utilities such as WinZip need to be decompressed prior to use.

153. Much of the audio and video content that was available on the Internet was encoded in standard formats such as MPEG for video and MP3 for audio. Appx. FF (Fallon, *The Internet Today*, 2001) at 89 (“MPEG ... is considered the de facto standard for motion video via the World Wide Web.”); Appx. GG (Underdahl, *Internet Bible*, 2000) at 292 (“Sources for MP3 music are cropping up everywhere, but almost all of them are online.”). Encoding media into these formats involves compression. *Id.* at 282 (The MPEG “format compresses sound and movie files to make them easier to move around the Internet.”); Appx. DD (Cadenhead, *How to Use the Internet*, 2000) at 143 (The MP3 “format was

developed with the goal of preserving sound quality while making files as small as possible.”). Proprietary encoding formats (e.g., Apple’s QuickTime) were also popular and also involved compression. *See, e.g., Appx. FF (Fallon, The Internet Today, 2001) at 88-89 (discussing QuickTime and AVI among other popular file formats on the Internet).*

154. To play media files, computers use applications to decode the format of the media file. Appx. GG (Underdahl, *Internet Bible*, 2000) at 283 (noting that player programs can play media in certain standard and proprietary formats). In addition to decoding, computers typically had to take digital audio and/or video and convert it to analog signals that could be output to analog devices (e.g., VGA monitors, TVs, powered speakers, headphones). This process is known as digital-to-analog conversion. The figure below depicts a generic digital video system dataflow where the decoder performs decoding then digital-to-analog conversion (“D to A”) to generate an analog video signal:

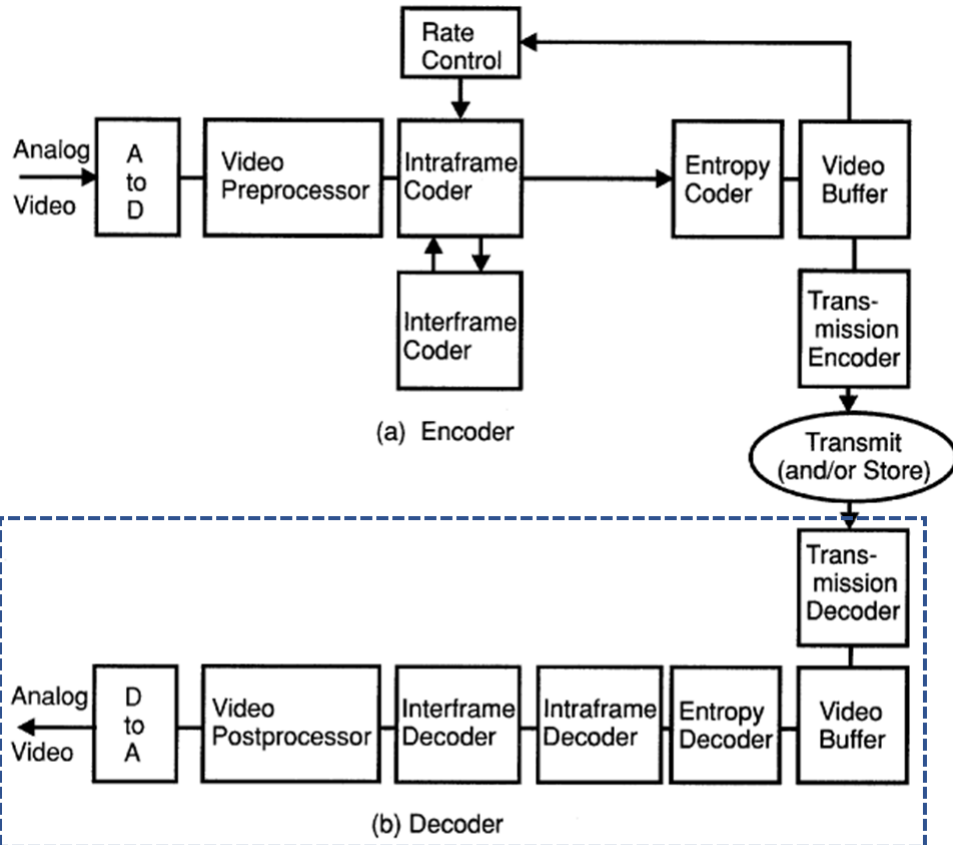


Figure 5.12 Digital video system dataflow.

Appx. HH (Hoffman, *Data Compression in Digital Systems*, 1997) at 106 (blue dashed box added). *See also*, e.g., Appx. II (Jurgen, *Digital Consumer Electronics Handbook*, 1997) at 25.4 (Figure 25.1 showing PC multimedia system with MPEG decoder and video and audio digital-to-analog decoders (DACs)); *see also id.* at 25.5 - 25.6 (discussing MPEG decoders).

155. The use of encryption and decryption in connection with data received over the Internet was also well-understood by late 2000. Secure Sockets Layer (SSL) could be used to encrypt communications between clients and servers on the Internet. Appx. Y (Rubin, *Web Security Sourcebook*, 1997) at 260 (“Netscape

Communications Secure Socket Layer (SSL) protocol is currently in its third revision and has been offered to the IETF as a candidate standard for HTTP security.”); *id.* at 263 (“SSL has garnered fairly strong support in the user community, primarily due to the large base of Navigator clients. Most commercial Web server software packages support SSL, and many free Web servers include hooks for integrating SSL.”); *see also* Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 711 (“You don’t have to do anything to get SSL; it is built into both Navigator and Internet Explorer.”). SSL allows users to exchange encrypted data after completing an initial handshake process. Appx. Y (Rubin, *Web Security Sourcebook*, 1997) at 262 (“Once SSL completes the handshake phase, it enters into an opaque data mode, in which application data is passed in encrypted, sequenced chunks, each including a cryptographic checksum to prevent tampering.”); Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 711 (“SSL provides browser and Web servers three important security services: encryption, certificate handling and message integrity.”).

156. It was also well-understood by late 2000 that files exchanged online could be protected with encryption prior to transfer. *See, e.g.*, Appx. S (Pike, *Using FTP*, 1995) at 23 (“[I]t makes sense to protect important information (by using encryption encoding, for example) before transferring it using FTP. There are several different commercial encryption products available to let you encrypt

and decrypt your data.”). A particularly well-known encryption utility was PGP. Appx. X (Ellsworth, *Internet 1997 Unleashed*, 1997) at 277-284 (describing PGP and how to use it). A file that was encrypted using a utility like PGP and then transferred would have to be decrypted prior to use.

(x) Step 10: Delivering Content to an Output Device

157. Element 1[o] requires delivering the processed output data via the wired connection to the output device:

(10) delivering, by the output system, via the wired connection, from the output system, and to the at least an output device wire connected to the output system, the audio output data or the video output data, the audio output data or the video output data is related to the output data wirelessly received from the one or more servers in (8) and that is processed from at least part of the output data in (9), and the audio output data or the video output data is for playing, at the output or at the at least an output device, at least part of the digital content selected in (6); and

158. The '903 patent does not describe any new or improved way of delivering output data via a wired connection from an output controller or any other device to an output device. The patent generally states that conventional “[w]ired links ... such as parallel interface, USB, Firewire interface, Ethernet and token ring networks may ... be implemented in the present invention” *Id.* at 20:50-53. Although the claims specifically require delivering “audio output data” or “video output data,” there is no particularized description of delivering such data

to an output device and no description of how an output controller or any other device is wire connected to any audio or video output device.

159. As I discussed above in discussing the claimed “output system,” there was nothing unconventional about delivering audio or video data over a wired connection to an output device. Conventional PCs included audio outputs for outputting audio signals to speakers and VGA ports for outputting video signals to monitors. Some computers even included S-video or composite video outputs for outputting video signals to TVs. *See supra* ¶ 96.

(xi) Step 11: Executing a Discovery Operation

160. Element 1[q] requires “executing a wireless discovery operation” to allow the output system to be discovered by a “client device”:

executing a wireless discovery operation, by the output system using the at least one chip or chipset of the output system, the execution of the wireless discovery operation is for a client device, which is in the same wireless local area network as the output system, to wirelessly discover, over the wireless local area network wirelessly coupled by the output system in (2), the output system for output service, the client device being a separate device from the output system and from the one or more servers; and

161. The '903 patent describes an optional discovery process that may be used by an information apparatus to find output devices. '903 patent at 22:19-28; 23:46-25:22; Fig. 5. The patent does not suggest that the invention requires any new or improved discovery process, however. On the contrary, the patent notes

that protocols such as Bluetooth, Salutation, Service Location Protocol (SLP), and Universal Plug and Play (UPnP) can be used. *Id.* at 23:66-24:9. Additionally, widely used operating systems such as Windows and Mac OS included functionality to allow users to discover printers available on a network. Appx H (Poole, *The Little Network Book for Windows and Macintosh*, 1999) at 133-145.

162. The use of discovery processes to allow devices to find one another was well-known by November 2000. The Bluetooth specification, published in 1999, describes a Service Discovery Protocol (SDP). Appx. JJ (*Bluetooth Core Specification*, 1999) at 323-384. The Bluetooth specification explains that the service discovery “problem is widely recognized; many companies, standards bodies and consortia are addressing it at various levels in various ways. Service Location Protocol (SLP), Jini, and Salutation, to name just a few, all address some aspect of service discovery.” *Id.* at 370. Service Location Protocol (SLP) was described in RFC 2165⁹, published in 1997. Appx. MM attached to this declaration is an IEEE Internet Computing article describing SLP. Appx. MM (Guttman, *Service Location Protocol: Automatic Discovery of IP Network Services*, 1999). UPnP relies on the Simple Service Discovery Protocol (SSDP) which was documented in several IETF drafts published in 1999. *See, e.g.*, Appx.

⁹ See <https://datatracker.ietf.org/doc/html/rfc2165>.

OO (Goland, *Simple Service Discovery Protocol/1.0*, Oct. 28, 1999). The Richard article I mentioned above describes the discovery capabilities in Bluetooth, Salutation, SLP, and UPnP. Appx. E (Richard, *Service Advertisement and Discovery: Enabling Universal Device Cooperation*, 2000).

(xii) Step 12: Receiving Content from a Client Device

163. Element 1[r] requires wirelessly receiving audio or video digital content from the client device:

wirelessly receiving, by the output system, using the at least one chip or chipset of the output system that is compatible, at least in part, with at least part of a protocol within IEEE 802.11 wireless standards, from the client device that has wirelessly discovered the output system, and over the wireless local area network wirelessly coupled by the output system in (2), audio or video digital content that includes audio data or video data for playing at the output system or at the at least an output device.

164. The '903 patent does not describe any new or improved way of wirelessly transmitting audio or video content to an output system. The description of how information apparatus 100 transmits output data to an output device is completely generic. *Id.* at 32:46-48 (“After receiving output data from server application 112, the information apparatus 100 may then transmit the output data to the output device 106 selected by the user in step 806.”).

165. By late 2000, there was nothing unconventional about transferring

audio or video data between computers on a wireless LAN. Audio and video files (e.g., MP3 files) could be transmitted like any other type of data file between computers on a wireless LAN. For example, one computer could set up a share and another could access that share and copy an audio or video file to or from it. *See, e.g.,* Appx. H (Poole, *The Little Network Book for Windows and Macintosh*, 1999) at 147-179 (discussing network file sharing in Windows and Mac OS); Appx. KK (Pogue, *Mac OS 9: The Missing Manual*, 2000) at 281-292 (discussing file sharing in Mac OS 9). An article in the June 2000 issue of MacWorld described wirelessly streaming MP3 files from one computer to an iBook in order to listen to music through home stereo speakers. Appx. LL (Breen, *Cut Loose*, June 2000) at 86 (“Rig Up a Remote Jukebox”).

166. FTP could also be used to exchange files on a local network. Windows NT Server included Internet Information Server (IIS) which had both FTP server and web server capabilities. *See, e.g.,* Appx. G (Russel, *Running Microsoft Windows NT Server 4.0*, 1997) at 402-410. The '903 patent itself states that a “server” may be used as an output controller. '903 patent at 14:66-15:4 (“Using a PC, **server**, workstation, or other computer to implement the feature sets of output controller 104 with application software is ***just another possible embodiment of the output controller 104*** and in no way departs from the spirit, scope and process of the present invention.”). These approaches work regardless

of whether either or both the sending and receiving computers are connected wirelessly to the network. In general, the use of wireless networking impacts low level networking protocols, not application layer protocols or user applications that rely on networking.

167. Audio or video data can be large and, as a result, can take longer to transfer over a network connection regardless of whether the network is wired or wireless. The size of audio and video data is impacted significantly by the length and quality of content and the compression used. An uncompressed WAV file can be encoded and compressed into a much smaller MP3 file, for example. The 802.11b extension published in 1999 added support for data rates up to 11 Mbps.¹⁰ Data rates of 5 Mbps were achievable with commercially available products. *See, e.g.,* Appx. M (Derfler, *Wireless LANs*, PC Magazine, April 18, 2000) at 236 (reporting throughput based on testing of several wireless LAN products). At a data rate of 5 Mbps, a 5 MB file (*e.g.,* an MP3 file) could be wirelessly transferred in 8 seconds. *See also, e.g.,* Appx. LL (Breen, *Cut Loose*, June 2000) at 86 (listing times between 37 and 77 seconds to transfer a 20 MB file using Apple's AirPort 802.11 technologies depending on the network configuration).

¹⁰ *See* Appx YY (IEEE Std. 802.11b-1999) at 11 (“This extension of the DSSS system builds on the data rate capabilities, as described in Clause 15 of IEEE Std. 802.11, 1999 Edition, to provide 5.5 Mbit/s and 11Mbit/s payload data rates in addition to the 1 Mbps and 2Mbps rates.”).

e. The Claim Elements as an Ordered Combination

168. In my opinion, there is nothing unconventional about the ordering or arrangements of the elements in claim 1. The “claimed” output system interacts with the “one or more servers” in ways that were routine and conventional. The output system obtains authentication information and uses it to access a service. The output system allows a user to select content and downloads that content from the server. The output system processes (*e.g.*, decompresses or decodes) the downloaded content and then delivers it to an output device. There is nothing unconventional about downloading data prior to processing it and delivering it to an output device. All of this is consistent with a PC downloading audio or video content from a secure web site or FTP server and then playing the content on their PC through a monitor and/or speakers.

169. The last two steps in claim 1 recited in elements 1[q] and 1[r] require the output system to be discovered by a “client device” and to receive audio or video from the client device. I see nothing unconventional about this ordering or arrangement either. The purpose of discovery was often to find other devices on a network to provide some service (*e.g.*, printing, projection). *See, e.g.*, Appx. E (Richard, *Service Advertisement and Discovery: Enabling Universal Device Cooperation*, 2000) at 18 (explaining that discovery can be used to help computers “establish *connections to neighboring devices for storage, faxing*, high-speed

network access, and **printing**"); Appx. MM (Guttman, *Service Location Protocol: Automatic Discovery of IP Network Services*, 1999) at 74 ("In this example, a client program seeks ***an overhead projection server to display a presentation*** to an assembled audience.").

170. The last two steps in claim 1 are generally unrelated to the ten enumerated steps that make up the bulk of the claim. There is no relationship required between the output data received from the client device and the output data received from the server. In my opinion, there was nothing unconventional about a computing device that ***both*** (1) downloads content from a server and also (2) receives content from a client device. By late 2000, conventional PCs could download content from the Internet and also receive content from other computers on a wired or wireless LAN. Moreover, as I discussed above, the '903 patent actually describes a system in which the information apparatus (or client device) downloads the content and then relays it to the output device. All of the output data transmitted to the output system comes through the same path—from the application server 110 to the information apparatus 100 and then to the output device. There is no discussion of any device that separately receives content from a server and from a client device, nor any suggestion that receiving content from those two different sources requires any new or improved technology.

2. Independent Claims 8 and 15

171. Independent claims 8 and 15 are similar to claim 1 discussed above. In my opinion, these claims also recite computer technology that was routine and conventional by late 2000. In this section, I will address the additional features recited in these claims that are not recited in claim 1.

172. Claim 8 uses the term “wireless output device” instead of “output system.” The term “wireless output device” is not used in the specification. Similar to the “output system” of claims 1 and 15, the claims themselves confirm that the “wireless output device” of claim 8 can be “a controller box connected to a television.” Ex. 1003, claim 13. As I discussed above, the ’903 patent describes the output controller in functional terms and explains that it can be implemented with a *conventional PC*. ’903 patent at 14:59-15:4. Thus, the specification confirms that the “wireless output device” of claim 8 can be a conventional PC.

173. Claim 8 requires the “one or more wireless communication units” to include “one or more radio frequency link controllers.” The ’903 patent describes an RF Link Controller 210 only briefly and in functional terms. ’903 patent at 19:48-55 (“RF link controller 210 implements real-time lower layer (e.g., physical layer) protocol processing that enables the hosts (e.g., information apparatus 100, output controller 104, output device 106, etc.) to communicate over a radio link. Functions performed by the link controller 210 may include, without limitation,

error detection/correction, power control, data packet processing, data encryption/decryption and other data processing functions.”). The ’903 patent specification thus reveals that the claimed “link controller” is a functional block within the wireless communication unit that enables conventional wireless communication. *See* ’903 patent at 20:45-49 (noting that a wireless unit like that shown in Fig. 2A “may be included in devices ... to support various wireless communications standards”).

174. Whereas the job object in claim 1 is required to include “authentication information,” the job object in claim 8 can include “authentication information, payment information, or subscription information, individually or in combination.” Ex. 1003, element 8[g]. Claim 15 recites that the job object contains authentication and/or subscription information. *Id.*, element 15[h]. The ’903 patent’s description of submitting payment information describes nothing new. ’903 patent at 34:47-48 (“If payment by credit card is selected, the user may be prompted to provide credit card information.”). By late 2000, it was routine and conventional to submit credit card information to a server on the Internet. *See, e.g.,* Appx. NN (Brown, *Using Netscape Communicator 4*, 1997) at 170-177 (describing online purchase process using Amazon.com). The ’903 patent does not clearly differentiate between “subscription information” and “authentication information.” The patent states that “[i]f the output service is a subscription, then

membership, login, authorization, or security information may be collected instead or in addition [to payment information].” ’903 patent at 34:67-35:3. As I have discussed, it was commonplace for clients to submit username (login) and password (authorization) information in order to access web servers and FTP servers.

175. Element 8[m] requires the output device “to store, to queue, or to spool ... at least part of the output data” received from the server. There was nothing unconventional about storing data received from a server. Files downloaded from FTP sites and web sites were stored by the client computer. *See* above discussion of step 8 of claim 1.

176. Elements 8[o] and 15[p] are similar to element 1[q] but more specifically require the wireless output device / output system to broadcast its availability. This functionality is only described in passing, and in functional terms, in the ’903 patent. ’903 patent at 24: 40-43 (“Alternatively or in combination, an information apparatus 100 may “listen to” service broadcasts from one or more output devices 106 and then identify the one or more output devices 106 that are needed or acceptable.”). The patent suggests that known protocols such as Universal Plug and Play (UPnP) can be used for discovery. *Id.* at 24:2-5. UPnP’s discovery framework, Simple Service Discovery Protocol (SSDP), provides for service providers to send broadcasts to announce their presence.

Appx. OO (IETF Draft, *Simple Service Discovery Protocol/1.0*, Oct. 28, 1999) at 4 (“SSDP services may send HTTP UDP notification announcements to the SSDP multicast channel/port to announce their presence.”); *id.* (“[T]wo types of SSDP requests will be sent across the SSDP multicast channel/port.... The second are presence announcements, a SSDP service announcing its presence.”). Other discovery protocols included similar broadcast functionality. *See, e.g.,* Appx. MM (Guttman, *Service Location Protocol: Automatic Discovery of IP Network Services*, 1999) at 77 (“A variety of simple multicast discovery protocols have been proposed over the years.... Some of the proposals allow services to announce their presence as they come up and periodically thereafter, so clients can become immediately aware of new services.”).

177. Element 15[q] also requires the output system to transmit an identification attribute to the client device. The ’903 patent mentions attributes that provide “[i]dentification of an output device (e.g., brand, model, registration, IP address etc.).” ’903 patent at 6:63-64. There was nothing unconventional about devices exchanging identification information during discovery. For example, SSDP presence announcements broadcast by service provider devices include a USN (i.e., “a URI that uniquely identifies a particular instance of a service”), a service type identifier, and a location for contacting the service:

USN URI	Service Type URI	Expiration	Location
upnp:uuid:k91...	upnp:clockradio	3 days	http://foo.com/cr
uuid:x7z...	ms:wince	1 week	http://msce/win

Appx. OO (IETF Draft, *Simple Service Discovery Protocol/1.0*, Oct. 28, 1999) at 5. Other discovery frameworks such as SLP also provide for transmitting identification attributes. Appx. MM (Guttman, *Service Location Protocol: Automatic Discovery of IP Network Services*, 1999) at 73 (“Services are advertised using a Service URL, which contains the service’s location: the IP address, port number and, depending on the service type, path. Client applications that obtain this URL have all the information they need to connect to the advertised service.”). Additionally, the patent gives an “IP address” as an example of an identification attribute for an output device. IP addresses are necessarily transmitted whenever IP is used; IP packet headers include source and destination IP addresses.

178. Claims 8 and 15 also require “operating system software to facilitate download and installation of application software or software components at the [wireless output device / output system].” Ex. 1003, elements 8[a] and 15[a]. This feature is very briefly described in functional terms in the specification. ’903 patent at 15:21-23 (“Additional application software may be installed or upgraded to newer versions in order to, for example, provide additional functionalities or bug fixes.”); *see also* 16:38-43 (a configuration application on a host computer may be

used to download and install updated software to output controller 104). The discussion does not mention any new or specific technology or approach to downloading and installing software updates. Conventional operating systems (e.g., Windows, Mac OS) have long allowed users to download and install new software. Wireless access points included similar functionality. For example, Cisco Aironet 340 series access allowed firmware updates to be downloaded from a computer used to configure the access point or from a file server.

Fully Update Firmware

These are links to alternative ways for reading and updating system firmware, radio firmware and web pages, all in one step.

Through Browser

This method allows you to browse your hard drive or mapped network drives to find the desired firmware and web page files.

From File Server

With this method, you enter named file information for update from a file server.

Selectively Update Firmware

These are links to alternative ways for reading and updating system firmware, radio firmware and web pages. You can select which firmware to update (system firmware, radio firmware, or web pages) rather than updating them all at once.

Through Browser

This method allows you to browse your hard drive or mapped network drives to find the desired firmware and web page files.

From File Server

With this method, you enter named file information for update from a file server.

Appx. PP (*Using the Cisco Aironet 340 Series Access Point*, 2000) at 3-40 - 3-41.

Print servers (a type of output controller, *see* '903 patent at 14:59-61) included these capabilities as well. As an example, the Axis 540/640 Print Servers allowed firmware updates to be installed using FTP. *See, e.g.,* Appx. QQ (*Axis 540/640 User's Manual*, 1997) at 115-118.

3. Dependent Claims

179. In general, the dependent claims add limitations that narrow the independent claims in ways that were well-understood, routine and conventional by late 2000.

(i) Claims 2, 13 and 16

180. Dependent claims 2, 13, and 16 restrict the claimed “output system” / “wireless output device” and “client device” to certain types of devices. *See* Ex. 1003. Claims 2 and 13 require the client device to be a smart phone or information pad with a touch screen. In addition to these two devices, claim 16 permits the client device to be an Internet appliance or digital camera. All of these claims also require the “output system” / “wireless output device” to be at least one of an audio output device, a speaker, a projection device, a television, or a controller box connected to a television.

181. For the most part, the '903 patent describes a generic information apparatus 100. The '903 patent does not describe any new or improved smart phone, information pad, Internet appliance, or digital camera. These devices are only mentioned as possible examples of an information apparatus. *See, e.g.,* '903 patent at 1:61-2:4; 9:10-18 (both listing examples of information apparatuses). Digital cameras were well-known before 2000. *See, e.g.,* Appx. HH (Hoffman, *Data Compression in Digital Systems*, 1997) at 257-262 (discussing digital

cameras). In 2000, terms like “smart phone,” “information pad,” and “Internet appliance” were loosely defined and not always used consistently. In my opinion, known devices like the Nokia Communicator would have been viewed by a POSA as “smart phones.” *See, e.g.,* Appx. RR (Helal, *Any Time, Anywhere Computing: Mobile Computing Concepts and Technology*, 1999) at 40-41 (discussing Nokia Communicator). Furthermore, given that the '903 patent identifies “Internet appliances” as examples of “small and low-cost mobile device[]” (*see* '903 patent at 4:13-16) the Nokia Communicators would have been viewed as “Internet appliances” as well. The term “information pad” may refer to a device like the InfoPad developed at UC Berkeley:



Fig. 2. The InfoPad portable multimedia terminal.

Appx. SS (Truman, *The InfoPad Multimedia Terminal: A Portable Device for Wireless Information Access*, 1998) at 1075. The '903 patent does not describe

these types of devices or describe how or whether the alleged invention works any differently when these types of devices are used as opposed to other types of information apparatuses such as desktop and laptop computers.

182. The specific output devices recited in these claims are not meaningfully described in the '903 patent either. Like the client devices just discussed, the recited output devices are merely mentioned. *See, e.g.*, '903 patent at 2:4-7. The only output devices that receive any special attention in the specification are printers. *Id.* at 2:8-17; 12:21-13:43; Figs. 10A & 10B. Certainly televisions, projectors, and speakers were well-known by late 2000. Set-top boxes, DVRs, and digital satellite system receivers were well-known examples of external controller boxes wire connected to televisions. *See, e.g.*, Appx. HH (Hoffman, *Data Compression in Digital Systems*, 1997) at 250-252 (discussing digital set-top boxes); Appx. TT (Miller, *The Complete Idiot's Guide to Home Theater Systems*, 2000) at 109-114 (discussing DVRs), 115-122 (discussing digital satellite systems include set-top receivers). The '903 patent itself describes no particular controller box that is wire connected to a television; it only describes generic output controllers and output controllers for use with printers.

183. Accordingly, claims 2, 13, and 16 add nothing that was not already well-understood, routine, and conventional in my opinion.

(ii) Claims 3, 14, and 18

184. Dependent claims 3, 14, and 18 recite that the “output system” (or “wireless output device”) is a “controller box” that delivers audio or video output data to a television via a wired connection. Claim 14, like claims 2 and 13 discussed above, adds that the client device is at least a smart phone or information pad with a touch screen.

185. As I just noted, the '903 patent does not specifically describe any “controller box” or output controller for a television. A television is merely mentioned as a possible output device. *See, e.g.,* '903 patent at 2:4-7; 12:1-3. There is no specific discussion of processing audio or video data, or of outputting audio or video data via any particular type of wired connection to a television. Additionally, set-top boxes, DVRs, and digital satellite system receivers were well-known examples of external boxes wire connected to televisions. *See, e.g.,* Appx. HH (Hoffman, *Data Compression in Digital Systems*, 1997) at 250-252 (discussing digital set-top boxes); Appx. TT (Miller, *The Complete Idiot's Guide to Home Theater Systems*, 2000) at 109-114 (discussing DVRs), 115-122 (discussing digital satellite systems include set-top receivers).

186. Accordingly, claims 3, 14, and 18 add nothing that was not already well-understood, routine, and conventional in my opinion.

(iii) Claims 4, 11, and 19

187. Claims 4 and 11 require the output system / wireless output device to be capable of “direct short range wireless communication.” The output system / wireless output device must discover or be discovered by a wireless device, establish a direct short range wireless communication with the wireless device, and receive audio from the wireless device. Claim 19 is similar but specifically requires the use of Bluetooth.

188. Once again, none of this subject matter is described in any meaningful detail in the '903 patent. The patent briefly discusses Bluetooth and notes that it may be used as a wireless technology in practicing the invention. *Id.* at 19:56-20:19. There is no description of transmitting audio data in particular over Bluetooth or any other direct short range wireless communication.

189. Regardless, it was well understood by late 2000 that Bluetooth supported direct short range wireless communications and that it could be used for audio transmission. *See, e.g.,* Appx. UU (Haartsen, *Bluetooth—The Universal Radio Interface for Ad Hoc, Wireless Connectivity*, 1998) at 111 (discussing user scenarios include “The ultimate headset” and “Portable PC speakerphone”); Appx. JJ (*Bluetooth Core Specification*, v1.0B, 1999) at 985-992 (Appendix V: Bluetooth Audio). It was also well-known that IEEE 802.11 included an ad-hoc mode for direct wireless communications. *See, e.g.,* Appx. B (Geier, *Wireless LANs*, 1999)

at 108-109 (discussing ad hoc wireless networks). A POSA would have known that audio data (like any other data) could be transmitted using an 802.11 ad hoc wireless network.

190. Accordingly, claims 4, 11, and 19 add nothing that was not already well-understood, routine, and conventional in my opinion.

(iv) Claim 5

191. Claim 5 requires the output system to include a digital camera and recites four additional steps: (i) capture digital video data with the digital camera, (ii) send at least part of the captured video data to the server, and (iii) receive output data related to the uploaded video data back from the server, and (iv) output or play at least part of the received output data.

192. Nothing in claim 5 requires anything other than routine and conventional computer technology. As I discussed above, the '903 patent describes an alleged invention wherein content is uploaded (or at least identified) to a server that processes that content into output data for a selected output device. *See supra* ¶ 26. There is no specific discussion of uploading, processing, or downloading *video* data anywhere in the patent. Uploading files over the Internet was routine and conventional. FTP has traditionally allowed users not only to download files but also to upload them. *See, e.g.*, Appx. S (Pike, *Using FTP*, 1995) at 104-106 (describing file uploads using WS_FTP). HTML support for file

uploading through the <input type=file> tag was documented in November 1995 (RFC 1867¹¹) and subsequently incorporated into the HTML standard.

193. There is nothing new or unconventional about the receiving and outputting steps either. The '903 patent does not describe any new or improved way of downloading or outputting video or suggest that the invention would work any differently when the output data received from the server consists of video.

194. Claim 5 does not expressly require the server to process the uploaded video data into the downloaded output data. Regardless, even if such processing were required, it would have been routine and conventional. Commercially available video production software like Adobe Premiere could be used to take recorded video and encode it into a variety of output formats (RealPlayer, Apple QuickTime, AVI). *See, e.g.,* Appx. VV (Doucette, *Digital Video for Dummies*, 1999) at 223-228. (describing encoding for streaming with RealPlayer), 222-223, 229-235 (describing encoding to AVI and QuickTime formats); *see also* Appx. WW (*Adobe Premiere v5.0 - At a Glance*, 1998) at 2 (“New support for Web movies—Supports MPEG decoding and animated GIF output. Also provides a free plug-in that saves movies in Real Networks Real Video Streaming Format for use in Web pages. Support for new or updated video system software—Supports Apple

¹¹ *See* <https://datatracker.ietf.org/doc/html/rfc1867>.

QuickTime 3.0 for Macintosh and Windows, Microsoft DirectShow 5.0 (playback only), and Active Streaming Format (ASF).”).

195. Accordingly, claim 5 adds nothing that was not already well-understood, routine, and conventional in my opinion.

(v) Claims 6, 7 and 20

196. Dependent claims 6, 7, and 20 restrict the “interface” of the “output system” (or “wireless output device”). Ex. 1003. Claim 20 requires it to be at least a touch screen interface, a voice activated command interface, or a wireless interface. Claim 6 recites the same options, plus a graphical user interface (GUI). Claim 7 specifically requires a voice activated command interface. None of these interfaces were unconventional in late 2000.

197. The ’903 patent does not meaningfully describe any “voice activated command interface” or “touch sensitive screen interface.” These are merely mentioned as possible features. ’903 patent at 33:59-62 (“The user can make the selection by using, for example, any of a keyboard, keypad, mouse, stylus, soft keys, push buttons, software command, *touch sensitive screen, voice-activated command*, among others.”). Touch screen interfaces were widely used on PDAs and other handheld computers by 2000. *See, e.g.,* Appx. B (Geier, *Wireless LANs*, 1999) at 31 (“Common features of handheld PCs include ... Embedded touch screen with resolutions of 480x240 or 640x240 pixels, four gray scales (2-bit pixel

depth).”). Mac OS included voice command functionality called “PlainTalk.”

Appx. KK (Pogue, *Mac OS 9: The Missing Manual*, 2000) at 393-394 (“PlainTalk is what’s known as a command-and-control program. It lets you open programs, trigger AppleScripts, and click menu items by speaking their names.”). The ’903 patent does not describe any new GUI technologies either. *See, e.g.*, ’903 patent at 11:34-35 (client application 102 may “provide a graphical user interface (GUI) in its host information apparatus 100 to interact with user”). The GUI screens shown in Fig. 9 are simple, routine, and conventional screens that were typical of applications running on PDAs and handheld computers in 2000. Conventional PCs, laptops, and handheld computers in 2000 featured GUI-based operating systems (Windows, Mac OS) and GUI-based applications (web browsers, word processing applications). As I have discussed already, the ’903 patent does not describe any new wireless interface technologies either. The patent suggests using standard wireless technologies (*e.g.*, 802.11, Bluetooth). ’903 patent at 9:57-67.

198. Accordingly, claims 6, 7, and 20 add nothing that was not already well-understood, routine, and conventional in my opinion.

(vi) Claim 9

199. Dependent claim 9 recites that the authentication, payment, or subscription information included in the job object is obtained from the user through the interface of the wireless output device. *See* Ex. 1003. It also recites,

like claim 20, that the interface is at least a touch screen interface, a voice activated command interface, or a wireless interface.

200. As discussed above in addressing step 1 of claim 1, the '903 patent generically describes a user entering authentication information via a GUI. '903 patent at 27:14-16; 29:20-22; 31:48-52. Entry of payment information and subscription information are similarly mentioned without any technical discussion whatsoever. *Id.* at 34:47-48 (“If payment by credit card is selected, the user may be prompted to provide credit card information.”); 34:67-35:3 (“If the output service is a subscription, then membership, login, authorization, or security information may be collected instead or in addition.”).

201. As I discussed above in addressing step 1 of claim 1, entry of username and password information was routine and conventional. The '903 patent suggests that this same kind of information is subscription information. *Id.* at 34:67-35:3 (subscription services may collect “membership, login, authorization, or security information”). Entry of credit card information (“payment information”) was also conventional. Popular online shopping sites like Amazon.com allowed users to input credit card numbers to pay for purchases. *See, e.g.,* Appx. NN (Brown, *Using Netscape Communicator 4*, 1997) at 170-177 (describing use of Amazon.com including entry of credit card information to make purchases). The figure below shows an early version of the Amazon.com checkout

page which obtained an email address and password from the user (at least one of membership, login, authorization, or security information) and encouraged the user to submit their credit card information (payment information) online as well:

FIG. 10.12

Here's where you pay for
your purchase.

Amazon.com Order Form - page 1 - Netscape

File Edit View Go Communicator Help

Back Reload Home Search Guide Print Security

Bookmarks Netscape https://www.amazon.com/exec/obidos/order/2/6341-332670-331014

Amazon.com

Completing Your Order is Easy

We encourage you to enter your credit card number online (why this is safe). However, you also have the option of phoning us with the number after completing the order form. If you have any problems or questions, see the bottom of the page for details on our toll-free (800) customer support number.

1. Welcome.

Please enter your e-mail address: mbrown@ovalon.net

Please double check your e-mail address; one small typo and we won't be able to communicate with you about your order.

☐ I am a first-time customer. (You will be asked to create a password later on.)

☐ I am a returning customer, and my password is: [password field]

[Have you forgotten your password?](#)

2. Select a payment method.

Id. at 174 (highlighting added); *id.* (“There are several more pages of forms to fill out with credit card and shipping information before your order is completed, but you get the idea.”).

202. Finally, the interface limitation is routine and conventional, as I discussed for claims 6, 7, and 20.

203. Accordingly, claim 9 adds nothing that was not already well-understood, routine, and conventional in my opinion.

(vii) Claims 10 and 17

204. Dependent claim 10 requires the wireless output device to process the output data from the server by performing at least one of the same 16 generic processing operations recited in step (9) of claim 1. Claim 17 similarly recites that the processing of the output data in step (9) of claim 15 involves one of the same 16 generic processing operations. Claim 10 additionally recites that the “wireless output device” plays or outputs the audio or video output data after processing.

205. As I discussed above in addressing step (9) of claim 1, the recited processing operations were well-known, routine, and conventional, and they are only mentioned or discussed at a very high level in the '903 patent. As I have noted, there is no specific discussion in the '903 patent of processing audio or video data or how that processing differs at all from processing print data. Regardless, as I discussed above, conventional PCs were capable of decoding audio and video and converting digital audio and video to analog signals that could be output to monitors and speakers. *See supra* ¶ 153-154.

206. Conventional PCs could also output audio and video. *See, e.g.,* Appx. GG (Underdahl, *Internet Bible*, 2000) at 279-292 (describing playback of audio and video obtained from the Internet); Appx. XX (Allen, *The Learning Guide to the Internet*, 1997) at 123-138 (discussing playback of audio and video obtained via Internet Explorer). A desktop PC could output video through a monitor and

audio through speakers. A laptop PC could output audio and video through its built-in speaker(s) and display. As I have noted, some computers like Apple's PowerBook G3 series and the iBook models released in September 2000 included ports for outputting video signals to televisions. Conventional PCs could also be used to output audio to home stereo systems. *See, e.g.,* Appx. LL (Breen, *Cut Loose*, June 2000) at 86 ("Rig Up a Remote Jukebox").

207. Accordingly, claims 10 and 17 add nothing that was not already well-understood, routine, and conventional in my opinion.

(viii) Claim 12

208. Dependent claim 12 recites a generic "security procedure that includes security key authentication for establishing a restricted wireless communication link between the wireless output device and the client device." *See* Ex. 1003. The digital content received from the client device is received over the restricted wireless communication link" and then "render[ed] on a display screen or a television or a projection medium." *Id.*

209. The '903 patent mentions that an information apparatus may need to be authenticated in order to use a given output device. '903 patent at 27:45-47 ("With successful authentication, a user may gain access to all or part of the services provided by an output device 106"). The patent explains that "[a] simple authentication may be implemented by, for example, comparing the identity for the

information apparatus 100 with an approved control list of identifies or elements stored in the output device 106.” *Id.* at 27:34-37. The patent mentions that “[o]ther more complex authentication and encryption schemes may also be used,” but no such more complex authentication schemes are described. *Id.* at 27:37-39. “[S]ecurity keys (physical or digital)” are merely mentioned alongside seven other types of information that may be used to authenticate the user. *Id.* at 27:39-42 (“Information such as user name, password, ID number, signatures, **security keys (physical or digital)**, biometrics, fingerprints, voice, among others, may be used separately or in combination as authentication means.”). There is no specific discussion of how authentication is performed using security keys. There is also no specific discussion of how a “restricted wireless communication link” is established.

210. By late 2000, there was nothing unconventional about the claimed security procedure. PC users creating network shares could password protect them in order to restrict access. *See, e.g.,* Appx. H (Poole, *The Little Network Book for Windows and Macintosh*, 1999) at 147-179 (discussing network file sharing with access control in Windows and Mac OS); Appx. KK (Pogue, *Mac OS X: The Missing Manual*, 2000) at 286-291 (discussing access control in connection with shared disks and folders). Print servers, which the ’903 patent gives as examples of output controllers (’903 patent at 14:10-12), could also use security to restrict

access to shared printers. Appx. QQ (*Axis 540/640 User's Manual*, 1997) at 12 (“Main Features ... Security ... You may set up passwords for all users, restricting both login and printer access.”). To the extent a “restricted wireless communication link” would require the use of encryption,¹² the IEEE 802.11 standard provided for encryption. Appx. B (Geier, *Wireless LANs*, 1999) at 114 (discussing WEP).


211. Accordingly, claim 12 adds nothing that was not already well-understood, routine, and conventional in my opinion.

¹² See, e.g., '903 patent at 32:64-65 (describing transmission of output data from information apparatus 100 to output device 106: “Encryption techniques may be applied to the output data to ensure security.”).

IV. CONCLUSION

212. I declare under penalty of perjury under the laws of the United States that the foregoing is true and correct.

Executed on September 14, 2021 in Silver Spring, MD.

By: 
Samrat Bhattacharjee, Ph.D.

APPENDIX A

CURRICULUM VITAE BOBBY BHATTACHARJEE

DEPARTMENT OF COMPUTER SCIENCE
THE UNIVERSITY OF MARYLAND
COLLEGE PARK

1 Personal Information

Professor,
Computer Science Department and
the Institute for Advanced Computer Studies,
University of Maryland
Appointed Fall, 1999.

Affiliate Professor,
Department of Electrical and Computer Engineering,
University of Maryland.

Alfred P. Sloan Research Fellow (2004–2006).

1.1 Education

- Ph.D. in Computer Science
Georgia Institute of Technology, Atlanta, Georgia, Summer 1999
Dissertation title: Active Networking: Architectures, Composition, and Applications
Advisors: Kenneth L. Calvert and Ellen W. Zegura
- Bachelor of Science in Mathematics and Computer Science
Georgia College and State University, Milledgeville, Georgia, Spring 1994
Graduated *Summa Cum Laude* and Outstanding Department Major

1.2 Employment

Summer 2009 to present	Professor University of Maryland, College Park, Maryland
Summer 2005 to present	Associate Professor University of Maryland, College Park, Maryland
Fall 2006	Visiting Professor Max Planck Institut für Software Systems, Saarbrücken, Germany
Spring, Summer 2007	Visiting Researcher

APPENDIX A

	AT&T Labs, Florham Park, New Jersey
Fall 1999 to Spring 2005	Assistant Professor University of Maryland, College Park, Maryland
Fall 1995 to Summer 1999	Research Assistant Georgia Institute of Technology, Atlanta, Georgia
Summer 1998	Instructor Georgia Institute of Technology, Atlanta, Georgia
Summer 1997	Member of Technical Staff AT&T Labs, Florham Park, New Jersey
Summer 1995	Member of Technical Staff GTE Labs, Waltham, Massachusetts
Fall 1994 to Spring 1995	Teaching Assistant Georgia Institute of Technology, Atlanta, Georgia

2 Research, Scholarly, and Creative Activities

2.1 Chapters in Books

1. Gisli Hjálmtýsson and Samrat Bhattacharjee. “Control on Demand”, In *Proceedings of the First International Working Conference on Active Networks* volume 1653 of *Lecture Notes in Computer Science* (Stefan Covaci, editor), pages 315-329, Springer-Verlag, June 1999.
2. Pete Keleher, Samrat Bhattacharjee, and Bujor Silaghi. “Are Virtualized Overlay Networks Too Much of a Good Thing?”, *Peer-to-Peer Systems First International Workshop, Lecture Notes in Computer Science*, Vol. 2429, (Peter Druschel et. al. Editors) pages 225–231, Springer-Verlag, 2002.
3. Bobby Bhattacharjee, Sudarshan S. Chawathe, Vijay Gopalakrishnan, Peter J. Keleher, and Bujor D. Silaghi. “Efficient Peer-To-Peer Searches Using Result-Caching”, *Peer-to-Peer Systems II, Second International Workshop, IPTPS 2003, Lecture Notes in Computer Science*, Vol. 2735, (M. Frans Kaashoek and Ion Stoica, Editors), pages 225–236, Springer-Verlag, 2003.
4. Paolo Massa and Bobby Bhattacharjee. “Using Trust in Recommender Systems: an Experimental Analysis”, In *Second International Conference, iTrust 2004, Lecture Notes in Computer Science*, Vol. 2995 Jensen, Christian; Poslad, Stefan; Dimitrakos, Theo (Eds.), pages 221-235, Springer-Verlag, 2004.

5. Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. “Decentralized Message Ordering for Publish/Subscribe Systems”, *ACM/IFIP/Usenix 7th International Middleware Conference, Lecture Notes in Computer Science*, Vol. 4290 (Maarten van Steen and Michi Henning, Editors), pages 162–179, Springer-Verlag, 2006.
6. Misha Rabinovich and Bobby Bhattacharjee. “Overlay Networks and Resiliency”, in *Guide to Reliable Internet Services and Applications*, Charles R Kalmanek, Sudip Misra, and Y. Richard Yang (Editors). Springer-Verlag, 2010.

2.2 Articles in Refereed Journals

1. Kenneth L. Calvert, Samrat Bhattacharjee, Ellen W. Zegura, and James Sterbenz. “Directions in Active Networks”, *IEEE Communications Magazine*, No. 10, pages 72-78, 1998.
2. Samrat Bhattacharjee, Ellen W. Zegura, and Kenneth L. Calvert. “Active Networking and End-to-End Arguments”, *IEEE Network Magazine*, No. 3, pages 66-71, 1998.
3. Gisli Hjálmtýsson and Samrat Bhattacharjee. “Control on Demand - An Efficient Approach to Router Programmability”, *IEEE Journal on Selected Areas in Communications, JSAC*, Vol. 17, No. 9, pages 1549-1562, September 1999.
4. S. Bhattacharjee, W. C. Cheng, C.-F. Chou, L. Golubchik, and S. Khuller. “Bistro: a Platform for Building Scalable Wide-Area Upload Applications”, *ACM SIGMETRICS Performance Evaluation Review*, Vol. 28, No. 2, pages 29-35, September 2000.
5. Ellen W. Zegura, Mostafa Ammar, Zongming Fei, and Samrat Bhattacharjee. “Application-Layer Anycasting: A Server Selection Architecture and Use in Replicated Web Service”, *Transactions on Networking*, Vol. 8, Issue 4, pages 455-466, August 2000.
6. Suman Banerjee and Samrat Bhattacharjee. “Scalable Secure Group Communications over IP-multicast”, *IEEE Journal of Selected Areas in Communications, JSAC*, Vol. 20, No. 8, pages 1511 - 1527, October 2002.
7. U. Cetintemel, P. J. Keleher, B. Bhattacharjee, and M. J. Franklin. “Deno: A Decentralized, Peer-to-Peer Object-Replication System for Weakly-Connected Environments”, *IEEE Transactions on Computers*, Vol. 52, No. 7, pages 943–959, July 2003.
8. Suman Banerjee, Christopher Kommareddy, and Bobby Bhattacharjee. “Efficient Peer Location on the Internet”, *Computer Networks Journal*, Vol. 5:1, pages 5-17, 2004.
9. Rob Sherwood, Bobby Bhattacharjee, and Aravind Srinivasan. “P5: A Protocol for Scalable Anonymous Communications”, *Journal of Computer Security*, Vol 13:6, pages 839-876, 2005.
10. Suman Banerjee, Christopher Kommareddy, Koushik Kar, Bobby Bhattacharjee, and Samir Khuller. “OMNI: An Efficient Overlay Multicast Infrastructure for Real-time Applications”, *Special Issue of Computer Networks on Overlay Distribution Structures and their Applications*, Vol 50:6, pages 826-842, 2005.

APPENDIX A

11. Rob Sherwood, Seungjoon Lee, and Bobby Bhattacharjee. “Cooperative Peer Groups in NICE”, *Computer Networks Journal, Special Issue on Management in P2P systems: Trust, Reputation and Security*, Vol 50:4, pages 523-544, 2006.
12. Tuna Guven, Chris Kommareddy, Richard J. La, Mark A. Shayman, and Bobby Bhattacharjee. “Measurement-Based Optimal Routing on Overlay Architectures for Unicast Sessions”, *Computer Networks Journal: Special issue on Network Modeling and Simulation*, Vol. 50, No. 12, pages 1938–1951, August 2006.
13. Suman Banerjee, Seungjoon Lee, Bobby Bhattacharjee, and Aravind Srinivasan. “Resilient Multicast using Overlays”, *IEEE/ACM Transactions on Networking*, Vol. 14, No. 2, pages 237–248, April 2006.
14. Ruggero Morselli, Bobby Bhattacharjee, Michael A. Marsh, and Aravind Srinivasan. “Efficient Lookup on Unstructured Topologies”, *IEEE Journal on Selected Areas in Communications (J-SAC), Special Issue on Peer-to-Peer Communications and Applications*, pages 62-72, 2007.
15. Jik-Soo Kim, Beomseok Nam, Peter Keleher, Michael Marsh, Bobby Bhattacharjee, and Alan Sussman. “Trade-offs in Matching Jobs and Balancing Load for Distributed Desktop Grids”, *Future Generation Computer Systems – International Journal of Grid Computing: Theory, Methods & Applications*, Vol. 25, No. 5, pages 415–424, 2008.
16. Seungjoon Lee, Bobby Bhattacharjee, Aravind Srinivasan, and Samir Khuller. “Efficient and Resilient Backbones for Multihop Wireless Networks”, *IEEE Transactions on Mobile Computing*, Vol 7:11, 2008.
17. T. Guven, R. La, M. Shayman, and B. Bhattacharjee. “A Unified Framework for Multipath Routing for Unicast and Multicast”, *IEEE/ACM Transactions on Networking*, Vol. 6:5, 2008.
18. Seungjoon Lee, Bobby Bhattacharjee, Suman Banerjee, Bo Han. “A Generic Framework for Efficient Geographic Routing in Wireless Networks”, *Elsevier Computer Networks*, Vol 54:5, 2010.
19. Bo Han, Lusheng Ji, Seungjoon Lee, Bobby Bhattacharjee, and Robert R. Miller. “Are All Bits Equal? – Experimental Study of IEEE 802.11 Communication Bit Errors”. *IEEE/ACM Transactions on Networking*, Vol. 20, No. 6, 2012.
20. V. Singh, M. Lentz, B. Bhattacharjee, R. J. La and M. A. Shayman. “Dynamic frequency resource allocation in heterogeneous cellular networks”. *IEEE Trans. on Mobile Computing (TMC)*. 2016.
21. Wagner, Justin and Paulson, Joseph N. and Wang, Xiao and Bhattacharjee, Bobby and Bravo, Hector Corrada, *Privacy-Preserving Microbiome Analysis Using Secure Computation*, *Bioinformatics*, 2016.
22. Suman Banerjee, Bobby Bhattacharjee, and Christopher Kommareddy. “Scalable Application Layer Multicast”, *Transactions on Networking*, Under revision, Submitted in 2002.

2.3 Articles in Refereed Conferences and Workshops

1. Ellen W. Zegura, Kenneth L. Calvert, and Bobby Bhattacharjee. “How to Model an Internetwork”, In *Proceedings of INFOCOM’96*, pages 594–602, 1996.
2. Bobby Bhattacharjee, Mostafa Ammar, Ellen Zegura, Viren Shah, and Zongming Fei. “Application-Layer Anycasting”, In *Proceedings of INFOCOM’97*, pages 1388–1396, Kobe, Japan, 1997.
3. Bobby Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Active Networking and the End-to-End Argument”, In *Proceedings of ICNP’97*, pages 220–228, 1997.
4. Bobby Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “An Architecture for Active Networking”, In *Proceedings of IFIP TC6 Seventh International Conference on High Performance Networking’97*, pages 265–279, 1997.
5. Bobby Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Self-Organizing Wide Area Network Caches”, In *Proceedings of INFOCOM’98*, pages 600–608, 1998.
6. Bobby Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Reasoning about Active Networks”, In *Proceedings of ICNP’98*, pages 31–41, 1998.
7. Zongming Fei, Bobby Bhattacharjee, Ellen W. Zegura, and Mostafa Ammar. “A Novel Server Selection Technique for Improving the Response Time of a Replicated Service”, In *Proceedings of INFOCOM’98*, pages 783–791, San Francisco, CA, 1998.
8. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Congestion Control and Caching in CANEs”, In *Proceedings of ICC’98, Workshop on Active Networks and Programmable Networks*, 1998.
9. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “LIANE — Composition for Active Networks”, *Computer Communications Workshop*, 1998.
10. Gisli Hjálmtýsson and Samrat Bhattacharjee. “Control on Demand”, *Proceedings of International Workshop on Active Networking*, Berlin, pages 315–329, 1999.
11. S. Merugu, S. Bhattacharjee, E. Zegura, and K. Calvert. “Bowman: A Node OS for Active Networks”. *Proceedings of IEEE Infocom*, pages 1127–1136, 2000.
12. Y. Chae, S. Merugu, E. Zegura, and S. Bhattacharjee. “Exposing the Network: Support for Topology Sensitive Applications”, *Proceedings of IEEE OpenArch*, pages 65–74, 2000.
13. Samrat Bhattacharjee, William Cheng, Chen-Fu Chou, Leana Golubchik, and Samir Khuller. “Bistro: A Platform for Building Scalable Wide-Area Upload Applications”, *Proceedings of PAWS Workshop*, 2000.
14. R. Jaegar, S. Bhattacharjee, J. K. Hollingsworth, R. Duncan, T. Lavian, and F. Travostino. “Integrating Active Networking and Commercial-Grade Routing Platforms”, *Usenix 2000 Workshop on Intelligence at the Edge*, 2000.

APPENDIX A

15. Suman Banerjee and Bobby Bhattacharjee. “Scalable Group Communication over IP Multicast”, In *Proceedings of the International Conference on Network Protocols*, pages 261–271, 2001.
16. Narendar Shankar, Christopher Komareddy, and Bobby Bhattacharjee. “Finding Close Friends over the Internet”, In *Proceedings of the International Conference on Network Protocols*, pages 301–311, 2001.
17. Matt Sanders, Ken Calvert, Bobby Bhattacharjee, Stephen Zabele, Mark Keaton, and Ellen Zegura. “Active Reliable Multicast on CANEs: A Case Study”, *IEEE OpenArch*, pages 49–62, 2001.
18. Suman Banerjee and Samrat Bhattacharjee. “Scalable Application-Layer Multicast for Content Distribution”, *Computer Communications Workshop*, 2001.
19. Suman Banerjee, Bobby Bhattacharjee, and Christopher Komareddy. “Scalable Application-Layer Multicast”, *Proceedings of ACM SIGCOMM*, pages 205–217, 2002.
20. Bobby Bhattacharjee, Matt Sanders, Shashidhar Merugu, Ken Calvert, and Ellen Zegura. “CANEs: An Execution Environment for Composable Services”, In *DARPA Active Networks Conference and Exposition (DANCE 2002)*, pages 255–267, 2002.
21. Laura Bright, Samrat Bhattacharjee, and Louiqa Raschid. “Supporting Diverse Mobile Applications with Client Profiles”, In *Proceedings of ACM Workshop on Wireless Mobile Multimedia (WoWMoM)*, pages 88–95, 2002.
22. Suman Banerjee, Christopher Komareddy, and Bobby Bhattacharjee. “Scalable Peer-Finding on the Internet”, *Proceedings of Globecom 2002*, pages 2217–2221, 2002.
23. Rob Sherwood, Bobby Bhattacharjee, and Aravind Srinivasan. “P5: A Protocol for Scalable Anonymous Communications”, In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 58–70, 2002.
24. Pete Keleher, Samrat Bhattacharjee, and Bujor Silaghi. “Are Virtualized Overlay Networks Too Much of a Good Thing?”, *First International Workshop on Peer-to-Peer Systems (IPTPS’02)*, 2002.
25. Bujor Silaghi, Bobby Bhattacharjee, and Pete Keleher. “Routing in the TerraDir Directory Service”, *Proceedings of SPIE/ITCom 2002*, Vol. 4868-30, pages 42–53, 2002.
26. Suman Banerjee, Seungjoon Lee, Bobby Bhattacharjee, and Aravind Srinivasan. “Scalable Resilient Multicast”, *Proceedings of ACM SIGMETRICS*, pages 102–113, 2003.
27. K-T Kuo, S. Phuvoravan, B. Bhattacharjee, R. J. La, M. Shayman, and H. S. Chang. “On the Use of Flow Migration for Handling Short-term Overloads”, *IEEE Globecom*, pages 3108–3112, 2003.
28. Suman Banerjee, Christopher Komareddy, Koushik Kar, Bobby Bhattacharjee, and Samir Khuller. “Construction of an Efficient Overlay Multicast Infrastructure for Real-time Applications”, *Proceedings of IEEE Infocom*, pages 1521–1531, 2003.

APPENDIX A

29. Seungjoon Lee, Rob Sherwood, and Bobby Bhattacharjee. “Cooperative Peer Groups in NICE”, *Proceedings of IEEE Infocom*, pages 1272–1282, April 2003.
30. K-T Kuo, S. Phuvoravan, T. Guven, L. Sudarsan, H. S. Chang, S. Bhattacharjee, and M. A. Shayman. “Fast Timescale Control for MPLS Traffic Engineering”, *Proceedings of Globecom 2003*, pages 3108–3114, 2003.
31. Bobby Bhattacharjee, Sudarshan Chawathe, Vijay Gopalakrishnan, Pete Keleher, and Bujor Silaghi. “Efficient Peer-To-Peer Searches Using Result-Caching”, *Second International Workshop on Peer-to-Peer Systems (IPTPS’03)*, 2003.
32. Ruggero Morselli, Jonathan Katz, and Bobby Bhattacharjee. “A Game-Theoretic Framework for Analyzing Trust-Inference Protocols”, *In the Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems*, June 2004.
33. Bujor Silaghi, Pete Keleher, and Bobby Bhattacharjee. “Multi-Dimensional Quorum Sets for Read-Few Write-Many Replica Control Protocols”, *In Proceedings of the IEEE/ACM CCGRID, 4th Fourth International Workshop on Global and Peer-to-Peer Computing (GP2PC)*, April 2004.
34. Suman Banerjee, Seungjoon Lee, Ryan Braud, Bobby Bhattacharjee, and Aravind Srinivasan. “Scalable Resilient Media Streaming”, *In Proceedings of ACM NOSSDAV’04*, pages 4–9, 2004.
35. Seungjoon Lee, Suman Banerjee, and Bobby Bhattacharjee. “The Case for a Multi-hop Wireless Local Area Network”, *In Proceedings of IEEE Infocom*, pages 894–905, 2004.
36. T.Guven, C. Kommareddy, R.J. La, M.A. Shayman, and B. Bhattacharjee. “Measurement Based Optimal Multi-path Routing”, *In Proceedings of IEEE Infocom*, pages 187–196, March 2004.
37. Ruggero Morselli, Bobby Bhattacharjee, Jonathan Katz, and Pete Keleher. “Trust-Preserving Set Operations”, *In Proceedings of IEEE Infocom*, pages 2231–2241, March 2004.
38. Rob Sherwood, Ryan Braud, and Bobby Bhattacharjee. “Slurpie: A Cooperative Bulk Data Transfer Protocol”, *In Proceedings of IEEE Infocom*, pages 941–951, March 2004.
39. V. Gopalakrishnan, B. Silaghi, B. Bhattacharjee, and P. Keleher. “Adaptive Replication in Peer-to-Peer Systems”, *In Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 360–369, March 2004.
40. David Hovemeyer, Jeff Hollingsworth, and Bobby Bhattacharjee. “Running on the Bare Metal with GeekOS”, *In Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, pages 315–319, March 2004.
41. Bujor Silaghi, Vijay Gopalakrishnan, Bobby Bhattacharjee, and Pete Keleher. “Hierarchical Routing with Soft-State Replicas in TerraDir”, *In Proceedings of the 18th IPDPS Conference*, pages 48–57, April 2004.

APPENDIX A

42. Rob Sherwood, Bobby Bhattacharjee, and Ryan Braud. “Misbehaving TCP Receivers Can Cause Internet-Wide Congestion Collapse”, *Proceedings of Computer and Communications Security (CCS)*, pages 383–392, 2005.
43. Ruggero Morselli, Bobby Bhattacharjee, Michael A. Marsh, and Aravind Srinivasan. “Efficient Lookup on Unstructured Topologies”, *Principles of Distributed Computing*, pages 77–86, 2005.
44. Seungjoon Lee, Bobby Bhattacharjee, and Suman Banerjee. “Efficient Geographic Routing in Multihop Wireless Networks”, *ACM MobiHoc 2005*, pages 230–241, 2005.
45. T. Guven, R. J. La, M. Shayman, and B. Bhattacharjee. “Measurement-based Multipath Multicast,” *In IEEE Global Internet Symposium*, pages 2803–2808, 2005.
46. Vahid Tabatabaee, Bobby Bhattacharjee, Richard La, and Mark Shayman. “Differentiated Traffic Engineering for QoS Provisioning”, *In the Proceedings of INFOCOM’05*, pages 2349–2359, 2005.
47. J. S. Kim, B. Nam, P. Keleher, M. Marsh, B. Bhattacharjee, and A. Sussman. “Resource Discovery Techniques in Distributed Desktop Grid Environments”, *Proceedings of the 7th IEEE/ACM International Conference on Grid Computing - GRID 2006*, pages 9–16, September 2006. *Best paper award*.
48. Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. “Decentralized Message Ordering for Publish/Subscribe Systems”, *ACM/IFIP/Usenix 7th International Middleware Conference*, 2006.
49. Abhishek Kashyap, Samrat Bhattacharjee, Richard La, Mark Shayman, and Vahid Tabatabaee. “Single-Path Routing of Time-varying Traffic”, *In the Proceedings of Globecom*, 2006.
50. Vasile Gaburici, Peter Keleher, and Bobby Bhattacharjee. “File System Support for Collaboration in the Wide Area”, *Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 26–36, 2006.
51. J. H. Li, M. Yu, R. Levy, and B. Bhattacharjee. “A Scalable Key Management and Clustering Scheme for Ad Hoc Networks”, *Proceedings of InfoScale*, pages 28–38, Hong Kong, 2006.
52. Vijay Gopalakrishnan, Bobby Bhattacharjee, and Peter Keleher. “Distributing Google”, *In 2nd IEEE International Workshop on Networking Meets Databases (NetDB’06)*, pages 33–39, April 2006.
53. Dave Levin, Rob Sherwood, and Bobby Bhattacharjee. “Fair File Swarming with FOX”, *Proceedings of the International Workshop on Peer-to-Peer Systems (IPTPS’06)*, 2006.
54. Alan Mislove, Massimiliano Marcon, Krishna Gummadi, Peter Druschel, and Bobby Bhattacharjee. “Measurement and Analysis of Online Social Networks”, *In Proceedings of the ACM/Usenix Internet Measurement Conference (IMC 2007)*, pages 29–42, 2007.

APPENDIX A

55. Jik-Soo Kim, Peter Keleher, Michael Marsh, Bobby Bhattacharjee, and Alan Sussman. “Using Content-Addressable Networks for Load Balancing in Desktop Grids”, *In Sixteenth IEEE International Symposium on High-Performance Distributed Computing (HPDC)*, pages 189–198, 2007.
56. Vijay Gopalakrishnan, Ruggero Morselli, Bobby Bhattacharjee, Peter J. Keleher, and Aravind Srinivasan. “Distributed Ranked Search”, *14th Annual IEEE International Conference on High Performance Computing (HiPC)*, pages 7–20, 2007. *Best paper award*.
57. Animesh Nandi, Aditya Ganjam, Peter Druschel, T. S. Eugene Ng, Ion Stoica, Hui Zhang, and Bobby Bhattacharjee. “A Shared Control Plane for Overlay Multicast”, *Fourth Usenix Symposium on Networked Systems Design and Implementation (NSDI 2007)*, 2007.
58. Seungjoon Lee, Dave Levin, Vijay Gopalakrishnan, and Bobby Bhattacharjee. “Backbone Construction in Selfish Wireless Networks”, *Proceedings of SIGMETRICS*, pages 121–132, 2007.
59. Jik-Soo Kim, Beomseok Nam, Michael A. Marsh, Peter J. Keleher, Bobby Bhattacharjee, Derek Richardson, Dennis Wellnitz, and Alan Sussman. “Creating a Robust Desktop Grid using Peer-to-Peer Services”, *Proceedings of NSF Next Generation Software Workshop (NSFNGS)* (Appears with proceedings of IPDPS 2007), pages 1–7, 2007.
60. Vahid Tabatabaee, Abhishek Kashyap, Bobby Bhattacharjee, Richard La, and M. Shayman. “Robust Routing with Unknown Traffic Matrices”, *IEEE INFOCOM Minisymposiums*, pages 2336–2440, 2007.
61. Adam Bender, Neil Spring, Dave Levin, and Bobby Bhattacharjee. “Accountability as a Service”, *USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI)*, 2007.
62. Dave Levin, Adam Bender, Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. “Boycotting and Extorting Nodes in an Internetwork”, *Workshop on the Economics of Networked Systems and Incentive-Based Computing*, 2007.
63. B. Bhattacharjee, R. Rodrigues, and P. Kouznetsov. “Secure Lookup without (Constrained) Flooding”, *Workshop on Recent Advances on Intrusion-Tolerant Systems (WRAITS)*, 2007.
64. R. Rodrigues, P. Kouznetsov, and B. Bhattacharjee. “Large-Scale Byzantine Fault Tolerance: Safe but Not Always Live”, *In the Third Workshop on Hot Topics in System Dependability (HotDep’07)*, 2007.
65. Ruggero Morselli, Bobby Bhattacharjee, Jonathan Katz and Michael Marsh. “Exploiting Approximate Transitivity of Trust”, *In Fourth International Conference on Broadband Communications, Networks, and Systems, 2007 (BroadNets 2007)*, pages 515–524, 2007.
66. Jik-Soo Kim, Beomseok Nam, Michael Marsh, Peter Keleher, Bobby Bhattacharjee, and Alan Sussman. “Integrating Categorical Resource Types into a P2P Desktop Grid System”, *In Proceedings of the 9th IEEE/ACM International Conference on Grid Computing (Grid 2008)*, 2008.

APPENDIX A

67. Dave Levin, Katrina LaCurts, Neil Spring, and Bobby Bhattacharjee. “BitTorrent is an Auction: Analyzing and Improving BitTorrent’s Incentives”, *In Proceedings of Sigcomm*, 2008.
68. Alan Mislove, Hema Swetha Koppula, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. “Growth of the Flickr social network”, *In Proceedings of the 1st ACM SIGCOMM Workshop on Social Networks (WOSN’08)*, Seattle, WA, 2008.
69. Dave Levin, Randolph Baden, Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. “Motivating Participation in Internet Routing Overlays”, *In NetEcon 2008 (Workshop on the Economics of Networks, Systems, and Computation)*, 2008.
70. Bo Han, Lusheng Ji, Seungjoon Lee, Bobby Bhattacharjee, Robert R. Miller. “All Bits Are Not Equal – A Study of IEEE 802.11 Communication Bit Errors”, INFOCOM 2009.
71. Bo Han, Lusheng Ji, Seungjoon Lee, Robert R. Miller, Bobby Bhattacharjee. “Channel Access Throttling for Overlapping BSS Management”. IEEE ICC, June 2009.
72. Vijay Gopalakrishnan, Bobby Bhattacharjee, K. K. Ramakrishnan, Rittwik Jana, Divesh Srivastava. “CPM: Adaptive Video-on-Demand with Cooperative Peer Assists and Multicast.” IEEE INFOCOM 2009.
73. Bo Han, Lusheng Ji, Seungjoon Lee, Robert Miller, Bobby Bhattacharjee, “Channel Access Throttling for Improving WLAN QoS”, IEEE Secon, Rome, Italy, June 2009
74. Cristian Lumezanu and Randy Baden and Dave Levin and Neil Spring and Bobby Bhattacharjee, “Symbiotic Relationships in Internet Routing Overlays”, Usenix-NSDI, 2009.
75. Cristian Lumezanu and Randy Baden and Neil Spring and Bobby Bhattacharjee, “Triangle Inequality and Routing Policy Violations in the Internet”, Passive and Active Measurement Conference (PAM), 2009.
76. Randy Baden, Adam Bender, Neil Spring, Bobby Bhattacharjee, Daniel Starin, “Persona: An Online Social Network with User-Defined Privacy”, Proceedings of SIGCOMM, August 2009
77. Cristian Lumezanu, Randolph Baden, Neil Spring, Bobby Bhattacharjee “Triangle Inequality Variations in the Internet”, Proceedings of IMC, November 2009
78. Randy Baden, Neil Spring, Bobby Bhattacharjee “Identifying close friends on the Internet”, The Workshop on Hot Topics in Networks (ACM Hotnets), 2009
79. Adam Bender, Rob Sherwood, Derek Monner, Nate Goergen, Neil Spring, Bobby Bhattacharjee, “Fighting Spam with the NeighborhoodWatch DHT”, Proceedings of IEEE INFOCOM, April 2009
80. Cristian Lumezanu, Randolph Baden, Dave Levin, Neil Spring, Bobby Bhattacharjee, “Symbiotic Relationships in Internet Routing Overlays”, Proceedings of USENIX NSDI, April 2009

APPENDIX A

81. Animesh Nandi, Bobby Bhattacharjee, Peter Druschel, “What a mesh: understanding the design tradeoffs for streaming multicast”, Extended Abstract, ACM SIGMETRICS Performance Evaluation, 2009.
82. Bo Han, Aaron Schulman, Francesco Gringoli, Neil Spring, Bobby Bhattacharjee, Lorenzo Nava, Lusheng Ji, Seungjoon Lee, and Robert Miller, “Maranello: Practical partial packet recovery for 802.11”, Proceedings of USENIX NSDI 2010.
83. Cristian Lumezanu, Dave Levin, Bo Han, Neil Spring, and Bobby Bhattacharjee, “Don’t love thy nearest neighbor”, International Workshop on Peer-to-Peer Systems (IPTPS), April 2010.
84. Cristian Lumezanu, Katherine Guo, Neil Spring, Bobby Bhattacharjee, “The Effect of Packet Loss on Redundancy Elimination in Cellular Wireless Networks”, ACM Sigcomm Internet Measurement Conference (IMC), 2010.
85. Satinder Pal Singh, Randolph Baden, Choon Lee, Bobby Bhattacharjee, Richard J. La, Mark Shayman, “IP Geolocation in Metropolitan Areas”, Extended Abstract, Proceedings of ACM Sigmetrics, 2011.
86. Matthew Lentz, Dave Levin, Jason Castonguay, Neil Spring, Bobby Bhattacharjee, D-mystifying the D-root address change, IMC (International Measurement Conference), 2013.
87. Lentz, Matthew and Erdélyi, Viktor and Aditya, Paarijaat and Shi, Elaine and Druschel, Peter and Bhattacharjee, Bobby, ”SDDR: Light-Weight, Secure Mobile Encounters”, USENIX Security Symposium, 2014.
88. Aditya, Paarijaat and Erdélyi, Viktor and Lentz, Matthew and Shi, Elaine and Bhattacharjee, Bobby and Druschel, Peter, “EnCore: Private, Context-based Communication for Mobile Social Apps”, International Conference on Mobile Systems, Applications, and Services (MobiSys), 2014.
89. Aditya, Paarijaat and Bhattacharjee, Bobby and Druschel, Peter and Erdélyi, Viktor and Lentz, Matthew, Brave New World: Privacy Risks for Mobile Users, Workshop on Security and Privacy Aspects of Mobile Environments (SPME), 2014.
90. Dave Levin, Youndo Lee, Luke Valenta, Zhihao Li, Victoria Lai, Cristian Lumezanu, Neil Spring, Bobby Bhattacharjee “Alibi Routing” Proceedigs of ACM SIGCOMM, 2015.
91. Matthew Lentz, James Litton, Bobby Bhattacharjee “Drowsy Power Management” Proceedings of Symposium on Operating Systems Principles (SOSP), 2015.
92. Raul Herbster, Scott DellaTorre, Peter Druschel, Bobby Bhattacharjee. “Privacy Capsules: Preventing Information Leaks by Mobile Apps” Proceedings of Mobisys, 2016.
93. Paarijaat Aditya, Rijurekha Sen, Seong Joon Oh, Rodrigo Benenson, Bobby Bhattacharjee, Peter Druschel, Tongtong Wu, Mario Fritz, Bernt Schiele. “I-Pic: A Platform for Privacy-Compliant Image Capture” Proceedings of Mobisys, 2016.

94. Vaibhav Singh, Matthew Lentz, Bobby Bhattacharjee, Richard La, Mark Shayman.
95. ‘Dynamic Frequency Resource Allocation in Heterogeneous Cellular Networks’
96. roceedings of IEEE TMC 2016 (IEEE Transactions on Mobile Computing)
97. James Litton, Anjo Vahldiek-Oberwagner, Eslam Elnikety, Deepak Garg, Bobby Bhattacharjee, Peter Druschel “Light-weight Contexts: An OS Abstraction for Safety and Performance” Proceedings of OSDI, 2016.
98. Zhihao Li, Dave Levin, Neil Spring, Bobby Bhattacharjee profile imageBobby Bhattacharjee. “Internet anycast: performance, problems, and potential” Proceedings of SIGCOMM, 2018.
99. Matthew Lentz, Rijurekha Sen, Peter Druschel, Bobby Bhattacharjee. “SeCloak: ARM TrustZone-based Mobile Peripheral Control” Proceedings of Mobisys 2018 (Conference on Mobile Systems, Applications, and Services)
100. Viktor Erdelyi, Trung-Kien Le, Bobby Bhattacharjee, Peter Druschel, Nobutaka Ono. “Sonoloc: Scalable positioning of commodity mobile devices.” In Proceedings of the Sixteenth International Conference on Mobile Systems, Applications, and Services (MobiSys 2018).
101. Lillian Tsai, Roberta De Viti, Matthew Lentz, Stefan Saroiu, Peter Druschel, Bobby Bhattacharjee. “enClosure: Group Communication via Encounter Closures” Proceedings of Mobisys 2019
102. Composing Abstractions using the null-Kernel James Litton, Deepak Garg, Peter Druschel, Bobby Bhattacharjee HotOS, 2019
103. Sergi Delgado-Segura, Surya Bakshi, Cristina Perez-Sola, James Litton, Andrew Pachulski Andrew Miller, Bobby Bhattacharjee “TxProbe: Discovering Bitcoin’s Network Topology Using Orphan Transactions” Financial Crypto, 2019
104. Venkat Arun, Aniket Kate, Deepak Garg, Peter Druschel, Bobby Bhattacharjee. “Finding Safety in Numbers with Secure Allegation Escrows” NDSS 2020 (to appear)

2.4 Technical Reports and Invited Papers

1. Ellen W. Zegura, Kenneth. L. Calvert, and Samrat Bhattacharjee. “Tera-op networking: Local adaptation to congestion”, In *Gigabit Networking Workshop*, 1996.
2. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Network Support for Multicast Video Distribution”, Technical Report GIT-CC-98-16, College of Computing, Georgia Institute of Technology, 1996.
3. Samrat Bhattacharjee. “Self-Organizing Wide-Area Network Caches”, Technical Report GIT-CC-97-31, College of Computing, Georgia Institute of Technology, 1996.

APPENDIX A

4. Zongming Fei, Samrat Bhattacharjee, Ellen W. Zegura, and Mostafa H. Ammar. “A Novel Server Selection Technique for Improving the Response Time of a Replicated Service”, Technical Report GIT-CC-97-24, College of Computing, Georgia Institute of Technology, 1996.
5. Samrat Bhattacharjee, Mostafa Ammar, Ellen Zegura, Viren Shah, and Zongming Fei. “Application-Layer Anycasting”, Technical Report GIT-CC-96-25, College of Computing, Georgia Institute of Technology, 1996.
6. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “An Architecture for Active Networking”, Technical Report GIT-CC-96-20, College of Computing, Georgia Institute of Technology, 1996.
7. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “On Active Networking and Congestion”, Technical Report GIT-CC-96-02, College of Computing, Georgia Institute of Technology, 1996.
8. Samrat Bhattacharjee, Ellen W. Zegura, and Kenneth L. Calvert. “High Performance Web: An Application for Wide Area Caching”, In *Gigabit Networking Workshop*, 1997.
9. Samrat Bhattacharjee and Gísli Hjalmtýsson. “Control-on-Demand: A Flow Oriented Approach towards Active Networking”, *AT&T Labs Technical Memorandum*, 1997.
10. Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. “Improving the Quality of Best Effort Service”, Technical Report GIT-CC-98-31, College of Computing, Georgia Institute of Technology, 1998.
11. S. Bhattacharjee and M. W. McKinnon. “Performance of Application-Specific Buffering Schemes for Active Networks”, Technical Report GIT-CC-98-17, College of Computing, Georgia Institute of Technology, 1998.
12. S. Merugu, S. Bhattacharjee, Y. Chae, M. Sanders, K. Calvert, and E. Zegura. “Bowman and CANEs: Implementation of an Active Network”, *Proceedings of the Thirty-Seventh Annual Allerton Conference on Communication, Control and Computing*, Monticello, Illinois, September 1999 (invited paper).
13. Ugur Cetintemel, Peter J. Keleher, and Bobby Bhattacharjee. “A Security Infrastructure for Mobile Transactional Systems”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4171, 2000.
14. Peter J. Keleher, Bobby Bhattacharjee, Kuo Kuo-Tung, and Ugur Cetintemel. “A Security Infrastructure for Mobile Transactional Systems”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4077, 2000.
15. Suman Banerjee and Bobby Bhattacharjee. “Scalable Secure Group Communication over IP Multicast”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4252, 2001.

APPENDIX A

16. Suman Banerjee and Samrat Bhattacharjee. “Spatial Clustering for IP Multicast: Algorithms and an Application”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4177, 2001.
17. Suman Banerjee and Bobby Bhattacharjee. “Analysis of the NICE Application Layer Multicast Protocol”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4380, 2002.
18. K-T Kuo, S. Phuvoravan, T. Guven, L. Sudarsan, S. Bhattacharjee, and M. A. Shayman. “Fast Time Scale Control for MPLS Traffic Engineering”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4351, 2002.
19. Suman Banerjee, Bobby Bhattacharjee, and Christopher Kommareddy. “Scalable Application Layer Multicast”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4373, 2002.
20. Bobby Bhattacharjee, Pete Keleher, and Bujor Silaghi. “The Design of TerraDir”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4299, 2002.
21. Suman Banerjee, Bobby Bhattacharjee, and Srinivasan Parthasarathy. “A Protocol for Scalable Application Layer Multicast”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4278, 2002.
22. Vijay Gopalakrishnan, Bujor Silaghi, Bobby Bhattacharjee, and Pete Keleher. “Adaptive Replication in Peer-to-Peer Systems”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4515, 2003.
23. Seungjoon Lee, Suman Banerjee, and Bobby Bhattacharjee. “The Case for a Multi-hop Wireless Local Area Network”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4504, 2003.
24. Christopher Kommareddy, Tuna Guven, Bobby Bhattacharjee, Richard La, and Mark Shayman. “Intradomain Overlays: Architecture and Applications”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4501, 2003.
25. Tuna Guven, Chris Kommareddy, Richard J. La, Mark A. Shayman, and Bobby Bhattacharjee. “Measurement Based Optimal Multi-path Routing”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4500, 2003.
26. Ruggero Morselli, Bobby Bhattacharjee, Jonathan Katz, and Pete Keleher. “Trust-Preserving Set Operations”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4499, 2003.
27. Suman Banerjee, Ryan Braud, Seungjoon Lee, Bobby Bhattacharjee, and Aravind Srinivasan. “Scalable Resilient Media Streaming”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4482, 2003.
28. Bujor Silaghi, Pete Keleher, and Bobby Bhattacharjee. “Multi-dimensional Quorum Sets for Read-Few Write-Many Replica Control Protocols”, University of Maryland, Dept. of Computer Science Technical Report, CS-TR-4440, 2003.

APPENDIX A

29. Tuna Guven, Richard La, Mark Shayman, and Bobby Bhattacharjee. “Measurement-based Multicast on an Overlay Architecture”, University of Maryland, Dept. of Computer Technical Report, CS-TR-4603, 2004.
30. Vijay Gopalakrishnan, Bobby Bhattacharjee, Sudarshan Chawathe, and Pete Keleher. “Efficient Peer-to-Peer Namespace Searches”, University of Maryland, Dept. of Computer Technical Report, CS-TR-4568, 2004.
31. Rob Sherwood, Bobby Bhattacharjee, and Ryan Braud. “Misbehaving TCP Receivers Can Cause Internet-Wide Congestion Collapse”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4737, 2005.
32. Jik-Soo Kim, Bobby Bhattacharjee, Peter Keleher, and Alan Sussman. “Matching Jobs to Resources in Distributed Desktop Grid Environments”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4791, 2006.
33. Ruggero Morselli, Bobby Bhattacharjee, Jonathan Katz, and Michael Marsh. “Key-Chains: A Decentralized Public-Key Infrastructure”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4788, 2006.
34. Ruggero Morselli, Bobby Bhattacharjee, Michael Marsh, and Aravind Srinivasan. “Efficient Lookup on Unstructured Topologies”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4772, 2006.
35. Seungjoon Lee, Bobby Bhattacharjee, and Suman Banerjee. “Efficient Geographic Routing in Multihop Wireless Networks”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4625, 2006.
36. Vijay Gopalakrishnan, Ruggero Morselli, Bobby Bhattacharjee, Peter Keleher, and Aravind Srinivasan. “Ranking Search Results in Peer-to-Peer Systems”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4779, 2006.
37. Randy Baden, Adam Bender, Dave Levin, Rob Sherwood, Neil Spring, and Bobby Bhattacharjee. “A Secure DHT via the Pigeonhole Principle”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4884, 2007.
38. Jik-Soo Kim, Peter Keleher, Michael Marsh, Bobby Bhattacharjee, and Alan Sussman. “Using Content-Addressable Networks for Load Balancing in Desktop Grids”, University of Maryland, Dept. of Computer Science Technical Report CS-TR-4863, 2007.

2.5 Tutorials, Talks, Abstracts, and Other Professional Papers Presented

- *Privacy by Design*, NEC Labs, Princeton, November 2012.
- *Systems without Cooperation*, South China University of Technology, October 2008.
- *Systems without Cooperation*, Sichuan University, October 2008.
- *Systems without Cooperation*, University of Maryland, September 2008.

APPENDIX A

- *Decentralized Applications on the Internet*, University of Lisbon, Portugal, July 2006.
- *Decentralized Applications on the Internet*, Bell Labs, January 2006, Murray Hill, New Jersey.
- *An Overview of Decentralized Applications*, at the *Algorithms in Networking* Workshop, FSCCTS 2005, Hyderabad, India.
- *Security Architectures for Peer-to-Peer Applications*, at the Marconi Foundation Video P2P Conference, Columbia University, 2004, New York City, NY.
- *Replication and Search in Distributed Namespaces*, at IBM Research, 2004, Hawthorne, NY.
- Invited Panelist, *Network Security: How Good Does it Have to Be?* at IEEE INFOCOM, 2003, San Francisco, CA.
- *P5: A Protocol for Scalable Anonymous Communications*, at IEEE S&P, 2002, Oakland, CA.
- *Cooperative Peer Groups in NICE*, at IEEE INFOCOM, April 2003, San Francisco, CA.
- *Overlay and P2P Systems: Protocols, Applications, and Analysis*, Tutorial (with Dan Rubenstein) at Networking Group Communications (NGC '02), October 2002, Boston, MA.
- *Cooperative Peer Groups in NICE*, Invited talk at BBN Technologies, Cambridge, MA, October 2002.
- *Finding Close Friends over the Internet*, at the International Conference on Network Protocols (ICNP), 2001, Riverside, CA.
- *Adaptive Network Processing*, at the Washington University Gigabit Switch Seminar, Washington University at St. Louis, St. Louis, January, 2001.
- *Active Networks: A Possible Future for the Internet?*, Invited Talk to the Washington DC/Northern VA Chapter for the IEEE/Microwave Theory and Techniques Society, April 2000.
- Invited Panelist at Gigabit Networking Workshop, San Francisco, CA, 1998 and International Communications Conference, Atlanta, GA 1998.
- *LIANE - Composition for Active Networks*, at IEEE Computer Communications Workshop, September 1998, Oxford, MS.
- *Self-Organizing Wide Area Network Caches*, at IEEE INFOCOM 1998, San Francisco, CA.
- *Reasoning about Active Networks*, at ICNP 1998, Austin, TX.
- *Finding the Best Server within the Application-Layer Anycasting Architecture*, at IEEE INFOCOM 98, San Francisco, CA.

APPENDIX A

- *High Speed Web: An Application for Active Caching*. Presented at Gigabit Networking Workshop '97, March 1997, Kobe, Japan.
- *Application-Layer Anycasting*, at IEEE INFOCOM 1997, Kobe, Japan.
- *Active Networking and the End-to-End Argument*, at IEEE ICNP'97, Atlanta, GA.
- *Tera-Op Networking: Local Adaptation to Congestion*. Presented at Gigabit Networking Workshop '96, March 1996, San Francisco, CA.

2.6 Patents

1. *Scalable wide-area upload system and method*, Leana Golubchik, William C. Cheng, Samir Khuller, Samrat Bhattacharjee, and Cheng-Fu Chou. United States Patent # 7,181,623. Granted: February 20, 2007.
2. *Method for encoding frame data*, Lusheng Ji, Samrat Bhattacharjee, Bo Han, Seungjoon Lee, Robert Miller. United States Patent # 7,940,850. Granted: May 10, 2011.
3. *Detection of distributed denial of service attacks in autonomous system domains*, Chris Kommareddy, Samrat Bhattacharjee, Mark Shayman, Richard La. United States Patent # 8,397,284. Granted: March 12, 2013.

2.7 Contracts and Grants

1. "EAGER: Decomposing Operating Systems for Better Control over Policy and Privacy", *National Science Foundation*, PI,
2. "LTS - Securing Critical Networking Infrastructure: DNS Root Servers", *Department of Defense*, PI, (Co-PIs: Neil Spring and David Levin), 2014-2015, \$199,799.00
3. "Interference Management in Heterogeneous Networks", *Air Force Research Laboratory*, Co-PI, (PI: Mark Shayman), 2012-2013, \$358,054.
4. "University Partnership with the Laboratory for Telecommunications Science", *Department of Defense*, Co-PI, 2010-2013, (PI: Joseph Jaja) \$896,814.
5. "University Partnership with the Laboratory for Telecommunications Science", *Department of Defense*, Co-PI, 2010-2013, (PI: Joseph Jaja) \$496,228.
6. "Privacy Preserving Social Systems", *National Science Foundation*, PI, Co-PIs: Neil Spring, Jonathan Katz, 2010 – 2013, \$880,000.
7. "Greed Resistant Protocols", *National Science Foundation*, Co-PI, (PI: Neil Spring), 2009 – 2012, \$499,344.
8. "An Integrated Approach to Computing Capacity and Developing Efficient Cross-Layer Protocols for Wireless Networks", *National Science Foundation*, Co-PI, (Principal Investigator: Aravind Srinivasan), September 2006 – September 2009, \$365,000.

APPENDIX A

9. “A Postmodern Internetwork Architecture”, *National Science Foundation*, Co-PI, Principal Investigator: Neil T. Spring, September 2006 – September 2009, \$400,000.
10. “Robust Grid Computing using Peer-to-Peer Services”, *NASA*, Co-PI. Principal Investigator: Alan Sussman. Other co-PIs: P. Keleher, D. Richardson, February 2006 – February 2009, \$1,008,251.
11. “A Wide-Area Event Notification System for MENTER”, *Laboratory for Telecommunication Sciences, National Security Agency*, January 2001 – August 2008, \$625,000.
12. “Employing Peer-to-Peer Services for Robust Grid Computing”, Co-PI. Principal Investigator: Alan Sussman. Other co-PIs: P. Keleher, D. Richardson, September 2005 – August 2006, \$60,000.
13. “Resilient Storage and Querying in Decentralized Networks”, *National Science Foundation*, Principal Investigator (Co-PI: Aravind Srinivasan, Sudarshan Chawathe, Jonathan Katz, Michael Marsh), Fall 2004 – Fall 2007, \$720,000.
14. Alfred P. Sloan Jr. Fellowship, September 2004 – September 2007, \$40,000.
15. “Distributed Trust Computations for Decentralized Systems”, *National Science Foundation*, Principal Investigator (Co-PI: Jonathan Katz), Fall 2003 – Fall 2006, \$375,000.
16. “CAREER: Adaptive Network Processing”, *National Science Foundation CAREER Award*, Fall 2001 – Spring 2006, \$500,000.
17. “Decentralized Directories for the Internet”, *National Science Foundation*, Principal Investigator (Co-PI: P. Keleher), Fall 2001 – Spring 2004, \$710,000.
18. “Parametric Design of Embedded Real-Time Systems”, *National Science Foundation*, Principal Investigator, Summer 2002 – Summer 2003. (Original PI: Richard Gerber, Fall 1998 – Summer 2002), \$200,154.
19. *Washington University Gigabit Switch Kit*. NSF, Washington University at St. Louis, Fall 1999 (Equipment only).

2.8 Fellowships, Prizes and Awards

1. Department of Computer Science Faculty Award for Teaching Excellence, 2012.
2. Department of Computer Science Faculty Award for Teaching Excellence, 2008.
3. Best paper award, 14th Annual IEEE International Conference on High Performance Computing (HiPC), 2007; paper co-authored with Vijay Gopalakrishnan, Ruggero Morselli, Peter J. Keleher, and Aravind Srinivasan.
4. Best paper award, 7th IEEE/ACM Conference on Grid Computing, 2006; paper co-authored with Jiksoo Kim, Byomsuk Nam, Peter Keleher, Michael Marsh, and Alan Sussman.

5. Alfred P. Sloan Jr. Fellowship, 2004.
6. Department of Computer Science Faculty Award for Teaching Excellence, 2004.
7. NSF CAREER award, 2001.
8. Recipient of Distinguished Teaching Assistant award from College of Computing, Georgia Tech, Spring 1997.

2.9 Editorial Boards and Reviewing Activities for Learned Publications

Reviewer for

ACM/IEEE Transactions on Networking

IEEE Journal on Selected Areas in Communications

Computer Communications Journal (Special Issue on Network Security)

ACM Transactions on Computer Systems

Performance Evaluation Journal

Computer Communications Review

European Transactions on Telecommunications

IEEE Transactions on Parallel and Distributed Systems

ACM Transactions on Internet Technology

Virtually all conferences in Networking and Systems including SOSP, OSDI, SIGCOMM, Sigmetrics, INFOCOM, SCW, DISC (formerly WDAG), Global Internet Conference, Infocom, IC3N, ICDCS, ICNP, ICPP, ICS, OpenArch, and WWW.

2.10 Research Software

1. *Odyssey: An active networking platform.* This distribution includes complete source and documentation for the Bowman Node OS and the CANEs Execution Environment. Released on the Internet, Summer 1999.
2. *NICE protocol simulator.* This distribution includes source for simulators of the NICE multicast protocols and complete implementation of the NICE protocols for video multicast. Released on the Internet, 2002.
3. *Slurpie.* This distribution includes the entire source code for a file-swarming system. Released on the Internet, 2004.
4. *OptAck Random Segment skip patch.* This software fixes a protocol fault (for Linux kernel versions 2.4 and 2.6). The fault exists in all known versions of TCP. Released on the Internet, 2005.
5. *Local Minima Search (LMS).* LMS is a protocol for unstructured search using virtual namespaces in distributed environments. Released on the Internet, 2006.

APPENDIX A

6. *KeyChains PKI*. KeyChains is a web-of-trust public key distribution/discovery system; it is built based on LMS local minima search algorithm, and uses CODEX libraries (from Cornell). Released on the Internet, 2006.
7. *Distributed Grid Software* The Distributed Grid software implements a complete distributed job matching system. The software suite is currently being field tested by researchers in Astronomy, and is available upon request, 2007.
8. *Cryptographic library for Chit-based access*. Developed cryptographic library for “chit”-based security. Library is used for different chit-based applications, including a filesystem and a distributed calendar application. Code available upon request, 2007.
9. *CPM on-demand video service*. The CPM software includes a novel video server and associated client software (and other supporting code) for implementing Cooperative Peer-Assisted Multicasting. Co-implemented the full software suite at AT&T Research. Code available upon request, 2007.
10. *IBOBSP*. IBOBSP is an in-network platform targeted towards reducing latency in interactive applications (in particular, games). The software distribution includes the in-network server pieces, and several graphical test applications and games. Co-implemented the full IBOBSP suite at AT&T Research. Code available upon request, 2007.

3 Teaching

3.1 Teaching Awards and Other Special Recognition

1. Teaching Excellence Award for Faculty, Department of Computer Science, Spring 2008.
2. Teaching Excellence Award for Faculty, Department of Computer Science, Spring 2004.
3. Distinguished Teaching Assistant, College of Computing, Georgia Tech, Spring 1997.

3.2 Advising: Research Advisor

3.2.1 Undergraduate

- Sebastian Gomez, Fall 2010 – Spring 2011.
- Chris Heistand, Fall 2010 – Spring 2011.
- Robert Kiefer, Spring 2009 – Summer 2010.
- Anika Cartas, Summer 2008 – Spring 2009.
- Katrina LaCurts, Fall 2007 – Summer 200.
- David Renie, Spring – Fall 2004.
- Ryan Evans Braud, Graduated Spring 2004.
- Mentor for Joseph Barrett, Colin Dixon, Tianzhou Duan, Kevin Genson, Bryant McIver, Ben Roseman, as part of the University of Maryland GEMSTONE program. Project title: *Anonymous Communications*, 2002-2005.

3.2.2 Masters

- Randolph Baden, Spring 2008.
- Chunyuan Liao, Fall 2004.
- Vijay Gopalakrishnan, Spring 2003.
- Kuo-Tung Kuo, Spring 2003.
- Dave Hovemeyer, Fall 2001.
- Vaibhav Kumar (ECE), Spring 2001.
- William Shapiro, Spring 2000.

3.2.3 Doctoral (completed)

- Suman Banerjee, graduated Summer 2003. Current position: Assistant Professor at University of Wisconsin.
- Laura Bright, graduated Spring 2003 (co-advisor). Current position: Research Associate, Oregon Graduate Institute.
- Vijay Gopalakrishnan, graduated Summer 2006. Current position: MTS, AT&T Research.
- Seungjoon Lee, graduated Summer 2006. Current position: MTS, AT&T Research.
- Ruggero Morselli, graduated Summer 2006. Current position: MTS, Google Inc.
- Christopher Kommareddy, graduated Summer 2006. Current position: Researcher, Amazon, Inc.
- Rob Sherwood, graduated Summer 2008. Current position: MTS, Deutsche Telekom Labs.
- Adam Bender, graduated Fall 2010, Current Position: MTS, Google.
- Dave Levin, graduated Summer 2010, Current Position: Visiting Research Professor, University of Maryland.
- Randolph Baden, graduated Summer 2012, Current Position: MTS, LTS-NSA.

3.2.4 Doctoral (current)

- Matthew Lentz
- Yeongsam Park
- Kookjin Lee
- James Litton
- Also advised visiting Ph.D. student Paolo Massa (Univ. of Trento) during Winter 2003-2004

3.3 Advising: Ph.D. Committees

- Nikhil Swami, Expected July 2008.
- Arun Vasan, 2008.
- Stephen Birrer (Northwestern University), 2007.
- Tuna Guven (ECE), 2006.
- Wan, Yung Chun, 2005.

APPENDIX A

- Andrzej Kochut, 2005.
- Yoo Ah Kim, 2005.
- Mehdi Kalantari (ECE), 2005.
- Arunesh Mishra, 2005.
- Surapich Phuvoravan (ECE), 2003.
- Bujor Silaghi, 2003.
- Suman Banerjee, 2003.
- Laura Bright, 2003.
- Kaushik Kar (ECE), 2002.
- Sungjoon Ahn, 2001.
- Ugur Cetintemel, 2001.
- Gabriel Rivera, 2001.
- Cuneyt Akinlar, 2001.
- Jung-Min Kim, 2001.
- Kritchalach Thitikamol, 2000.
- Saswati Sarkar (ECE), 2000.
- Demet Aksoy, 2000.

4 Service

4.1 Professional

4.1.1 Unpaid Reviewing Activities for Agencies

1. *NSF workshop on Network Testbeds*, attended workshop and co-authored report, 2002. Report basis for new NSF program on research testbeds.
2. NSF Networking Research Panel, Fall 2000, Fall 2001, Spring 2002, Fall 2002, Spring 2003, Spring 2004, Fall 2005, Spring 2008.
3. DoE High Performance Networking Panel, Spring 2001.
4. Evaluator for Intel Science Talent Search, 2001, 2002, 2003, 2004, 2005, 2007.

4.1.2 Other non-University Panels and Positions

- Program Committee Member, W-PIN+Netecon, 2013.
- Program Committee Member, WWW, 2013.
- Program Co-Chair, IEEE ICNP 2012.
- Program Committee Member, IMC, 2011.
- Co-Chair, Internet Research Task Force (IRTF) Peer-to-Peer Research Group, 2003 – 2009.
- Member, Internet Research Steering Group (IRSG), 2003 – current.
- Co Program Committee Chair, NetEcon 2009 Workshop.
- Program Committee Member, INFOCOM, 2009.
- Program Committee Member, Sigmetrics, 2009.
- Program Committee Member, Area TPC Chair, ICNP, 2008.
- Program Committee Member, INFOCOM, 2008.
- Program Committee Member, LANMAN, 2008.
- Program Committee Member, Sigmetrics, 2008.
- Program Committee Member, Workshop on Social Network Systems, 2008.
- Program Committee Member, ACM SIGCOMM, 2007.
- Program Committee Member, NetDB, 2007.
- Program Committee Member, Sigmetrics, 2007.
- Program Committee Member, ICDCS (P2P track), 2007.
- Program Committee Member, ICNP (P2P track), 2007.
- Program Committee Member, Electronic Commerce (EC), 2007.
- Area Chair (Dependable and Trustworthy Computing), ICPADS, 2007.
- Program Committee Member, IEEE Consumer Communications and Networking Conference - Workshop on Peer-to-Peer Multicasting, 2007.
- Program Committee Member, IEEE INFOCOM, 2006.
- Associate Chair (P2P, Grids Track), Fifteenth International Conference on Computer Communications and Networks (IC3N), 2006.

APPENDIX A

- Associate Chair (P2P Track), The 26th International Conference on Distributed Computing Systems (ICDCS), 2006.
- Program Committee Member, Global Internet (GI), 2006.
- Program Committee Member, IWAN, 2006.
- Program Committee Member, ICNP, 2006.
- Program Committee Member, IEEE INFOCOM, 2006.
- Program Committee Member, The 25th International Conference on Distributed Computing Systems (ICDCS), 2005.
- Program Committee Member, International Workshop on Active Networking (IWAN), 2005.
- Program Committee Member, IEEE INFOCOM 2005.
- Program Committee Member, IPTPS, 2005.
- Program Committee Member, HICSS, 2005.
- Program Committee Member, ACM SIGCOMM, 2004.
- Program Committee Member, 6th International Workshop on Distributed Computing (IWDC), 2004.
- Program Committee Member, 24th International Conference on Distributed Computing Systems (ICDCS), 2004.
- Program Committee Member, IEEE Global Internet Conference, 2004.
- Program Committee Member, International Workshop on Active Networking (IWAN), 2004.
- Program Committee Co-Chair, OpenArch, 2003.
- Program Committee Member, IEEE International Conference on Network Protocols, 2003.
- Program Committee Member, IEEE OpenSig, 2003.
- Program Committee Member, International Workshop on Networked Group Communications (NGC), 2003.
- Program Committee Member, International Workshop on Active Networking (IWAN), 2003.
- Program Committee Member, IEEE Global Internet Conference, 2003.
- Program Committee Member, IEEE International Conference on Network Protocols, 2002.

APPENDIX A


- Program Committee Member, International Workshop on Networked Group Communications (NGC), 2002.
- Program Committee Member, IEEE Global Internet Conference, 2002.
- Program Committee Member, International Workshop on Active Networking (IWAN), 2002.
- Program Committee Member, IEEE International Conference on Network Protocols, 2001.
- Program Committee Member, Workshop on Performance and Architecture of Web Servers (PAWS), 2001.
- Publications Chair, Member of Organizing and Program Committee, IEEE/ACM Open-Arch, 2001.
- Program Committee member, IEEE Global Internet Conference, 2001.



MACMILLAN
TECHNICAL
PUBLISHING
U.S.A.

APPENDIX B

MACMILLAN NETWORK
ARCHITECTURE &
DEVELOPMENT SERIES



WIRELESS LANs

Implementing Interoperable Networks

Jim Geier

ROKU EXH. 1002

WIRELESS LANs

Implementing Interoperable Networks

Jim Geier holds B.S.E. and M.S.E. degrees in electrical engineering, with an emphasis in computer networks. He was an active member of the IEEE 802.11 Working Group, responsible for developing international standards for wireless LANs. Jim has served as chairman of the Institute of Electrical and Electronic Engineers (IEEE) Computer Society, Dayton Section, and chairman of the IEEE International Conference on Wireless LAN Implementation. Jim has 18 years of experience providing information system consultation to companies worldwide, and has instructed many courses internationally on topics such as wireless networking, software development, and project management. He is currently the director of Data Collection Solution Development at Monarch Marking Systems. Jim is also the author of the *Wireless Networking Handbook* (1996, New Riders Publishing) and *Network Reengineering* (1996, McGraw-Hill), as well as numerous articles in leading publications, such as *Byte* and *Network Magazine*.

The *Macmillan Network Architecture and Development Series* is a comprehensive set of guides that provide computing professionals with the unique insight of leading experts in today's networking technologies. Each volume explores a technology or set of technologies that is needed to build and maintain the optimal network environment for any particular organization or situation.



CATEGORY: Networking

Wireless local area networks can provide unique benefits to many organizations, but require specific support and tools for maintaining network integrity. Based on the most recent developments in the field, *Wireless LANs*, gives network engineers, designers, and architects vital information on how to plan, configure, and implement wireless networks, including

- Coverage of the implications of migrating from proprietary solutions to the 802.11 standard
- Explanation of critical issues, such as maximizing interoperability between existing and future system infrastructure
- Authoritative advice on how to address common problems, such as radio frequency interference
- Discussion on how to realize significant cost savings through wireless LAN implementation for data collection systems
- Case studies and implementation notes, which provide real-world insight into the best practices of deploying a wireless LAN

This book provides both a context for understanding how an enterprise can benefit from the application of wireless technology, and the proven tools for efficiently implementing a wireless LAN. Designers and implementors will learn the considerations that must be addressed at each stage of the process, and find authoritative information on

- Primary wireless LAN applications, such as barcode scanners, data collectors, and printers
- The features and functionality of the IEEE 802.11 standard
- The details of upgrading from existing 902MHz to 2.4GHz networks
- Selecting the type of spread spectrum (direct sequence or frequency hopping) that best fits the needs of their particular networking environment

\$40.00 USA / \$57.95 CAN



ISBN 1-57870-081-7



ROKU EXH. 1002

Wireless LANs

Implementing Interoperable Networks

Jim Geier

M
TP
MACMILLAN
TECHNICAL
PUBLISHING
U.S.A.

Wireless LANs: Implementing Interoperable Networks

Copyright © 1999 by Macmillan Technical Publishing

FIRST EDITION

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

International Standard Book Number: 1-57870-081-7

Library of Congress Catalog Card Number: 98-85498

2001 00 99 98 4 3 2 1

Interpretation of the printing code: The rightmost double-digit number is the year of the book's printing; the rightmost single-digit, the number of the book's printing. For example, the printing code 98-1 shows that the first printing of the book occurred in 1998.

Composed in Bergamo and MCPdigital by Macmillan Computer Publishing

Printed in the United States of America

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Macmillan Technical Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

This book is designed to provide information about wireless LAN technology. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an as-is basis. The authors and Macmillan Technical Publishing shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

Feedback Information

At Macmillan Technical Publishing, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us at networktech@mcp.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher

Jim LeValley

Executive Editor

Linda Ratts Engelman

Managing Editor

Caroline Roop

Acquisitions Editor

Karen Wachs

Development Editor

Thomas Curtin

Project Editor

Laura N. Williams

Copy Editor

Keith Cline

Indexer

Tim Wright

Proofreader

Julie Searls

Acquisitions Coordinator

Amy Lewis

Manufacturing Coordinator

Brook Farling

Book Designer

Gary Adair

Cover Designer

Sandra Schroeder

Production Team Supervisor

Tricia Flodder

Production

Eric S. Miller

Supermarket scanners and most diffused infrared wireless LANs satisfy Class I requirements, where there is no hazard under any circumstance. Class IV specifies devices, such as laser-scalpels, which can cause grave danger if the operator handles them improperly. Most long-range, laser-based wireless networks are rated as Class III devices, whereby someone could damage his eyes if looking directly at the laser beam. Therefore, care should be taken when orienting lasers between buildings.

The Components of a Wireless Network

Wireless networks perform similar functions as their wired ethernet and token ring counterparts. In general, networks perform the following functions to enable the transfer of information from source to destination:

1. The medium provides a bit pipe (path for data to flow) for the transmission of data.
2. Medium access techniques facilitate the sharing of a common medium.
3. Synchronization and error control mechanisms ensure that each link transfers the data intact.
4. Routing mechanisms move the data from the originating source to the intended destination.
5. Connectivity software interfaces an appliance, such as pen-based computer or bar code scanner, to application software hosted on a server.

A good way to depict these functions is to specify the network's architecture. This architecture describes the protocols, major hardware, and software elements that constitute the network. A network architecture, whether wireless or wired, may be viewed in two ways, physically and logically.

Physical Architecture of a Wireless Network

The physical components of a wireless network implement the Physical, Data Link, and Network Layer functions (see Figure 1.7) to satisfy the functionality needed within local, metropolitan, and wide areas. The following sections explain the various components of a wireless LAN.

End-User Appliances

As with any system, there needs to be a way for users to interface with applications and services. Whether the network is wireless or wired, an *end-user appliance* is an interface between the user and the network. Following are the classes of end-user devices that are most effective as appliances for wireless networks:

- Desktop workstations
- Laptop computers

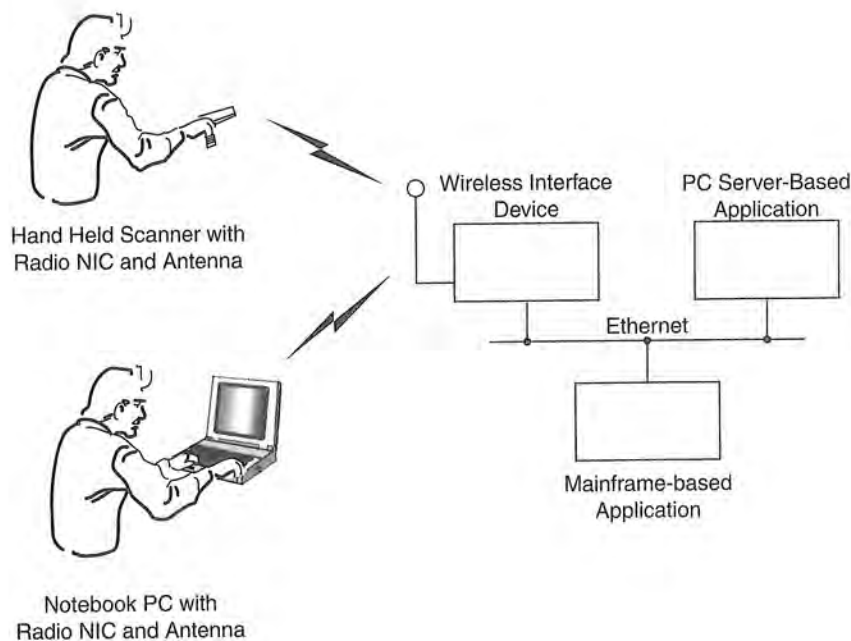


FIGURE 1.7 *The physical components of a wireless network extend the capability of ethernet and token ring.*

- Palmtop computers
- Handheld PCs
- Pen-based computers
- Personal digital assistants (PDA)
- Handheld scanners and data collectors
- Handheld printers

Today, the handheld PC, introduced by Microsoft (but developed and manufactured by other companies), is the primary hardware platform for Windows CE, which makes an excellent handheld wireless appliance. The main goals in developing the handheld PC include long battery life, affordable price (around \$500), compactness and light weight, familiar interfaces, easy PC connection, and effective keyboard input.

Note

Microsoft, being mostly a software house, signed up seven partners to develop a variety of handheld PCs that provide common functionality and vendor-specific features that support Windows CE. These partners are Casio, Compaq, Hewlett-Packard, Hitachi, Phillips Electronics, NEC, and LG Electronics.

Common features of handheld PCs include the following:

- Embedded QWERTY keyboard with alphanumeric keys, standard punctuation, a Ctrl key, an Alt key, and two Shift keys. Other vendor-specific keys are optional. A word of warning: If you have large fingers, you may have a difficult time pressing keys. Japanese and Chinese versions do not have keyboards; they have handwriting recognition as input.
- Embedded touch screen with resolutions of 480×240 or 640×240 pixels, four gray scales (2-bit pixel depth).
- Styles that acts like a mouse when tapped on the touch screen.
- Docking cradle to recharge the machine's batteries and connect it to your desktop PC.
- One PC Card (PCMCIA) slot, one serial connector, and one infrared port (IrDA).
- At least 2 MB RAM and 4 MB of ROM.

PalmPilot

As an example of handheld PCs, consider the PalmPilot by 3Com. It is a pocket-size organizer designed to connect seamlessly with a Windows-based or Macintosh computer. This combination of portability and one-touch connectivity provides a practical way to carry personal data anywhere. The PalmPilot fits in a shirt pocket and contains a suite of personal information management (PIM) applications.

A touchscreen and physical buttons provide one-finger data access. The compact Palm Connected Operating System switches screens and launches applications instantly, yet is efficient enough that two AAA batteries can power the device for several months. The organizer contains a memory module that the user can replace to add memory or upgrade the device. In addition, users will be able to attach communications add-on products,

such as modems and pagers as they become available.

The PalmPilot drops into a docking station that is connected to the desktop by a serial cable. Pressing the HotSync button on the cradle automatically backs up and synchronizes data with the desktop. Because the desktop synchronization software runs in the background, the user does not need to manage the process on the desktop and viewer. As a result, synchronizing data requires less user interaction than printing a document.

The PalmPilot includes Microsoft Windows or Macintosh OS companion versions of applications. Desktop software serves as the gateway between PalmPilot and desktop applications. For example, a mail merge between the PalmPilot Address Book and Microsoft

continues

continued

Word is accomplished with a simple click-and-drag operation.

Because wireless network appliances are often put into the hands of mobile people who work outdoors, the appliance must be tough enough to resist damage resulting from dropping, bumping, moisture, and heat. Some companies offer more durable versions of the portable

computer. Itronix, for example, sells the X-C 6000 Cross Country portable computer. The X-C 6000's case is built from strong, lightweight magnesium and includes a elastomer covering that protects the unit from weather and shock. The unit is impervious to rain, beverage spills, and other work environment hazards.

Note

When evaluating appliances for use with a mobile application, be certain to consider the ergonomics of the unit. You certainly won't be able to realize any of the benefits of a wireless network if users don't use the system because of appliances that weigh too much or are difficult to use.

Network Software

A wireless network consists of software that resides on different parts of the network. A network operating system (NOS), such as Microsoft NT Server, hosted on a high-end PC provides file, print, and application services. Many NOS's are server oriented, as shown in Figure 1.8, where the core application software and databases reside. In most cases, the appliances will interface via TCP/IP with application software or a database running on the NOS.

Client software, located on the end-user's appliance, directs the user's commands to the local appliance software, or steers them out through the wireless network. The software residing on a wireless appliance is very similar to software that runs on a wired appliance. The main difference is that it is important to develop the wireless software to optimize the use of the wireless network's relatively small amount of bandwidth.

The software performing application functions can run on a server/host, the appliance, or a combination of both. In some cases, such as with applications running on an IBM mainframe, IBM AS/400, or UNIX-based hosts, the wireless appliances may need to run terminal emulation. This makes the appliance act as a dumb terminal, just interfacing the keyboard, screen, printer, and so on, with the application running on the host. With client/server systems, the software on the appliance may perform part or all of the application's functionality and merely interface with a database located on a server, such as Microsoft NT Server. Chapter 6, "Wireless System Integration," covers this in more detail.




FIGURE 1.8 *The server-based network operating system provides a centralized platform for applications and data storage for mobile users.*

Note

A wireless network appears transparent to application software and operating systems on the network. As a result, applications written for a wired network can generally run without changes over a wireless network.

In some cases, a gateway running *middleware* is necessary to provide an interface between the appliance and the application software running on the server. The appliances communicate with the host/server through the gateway. The gateway acts as a proxy for the various appliances. The advantages of using the gateway are as follows:

- *Better RF throughput:* With the presence of a transport and application gateway, the appliances communicate with the gateway by using a “lightweight” protocol that is wireless friendly, unlike TCP/IP.
- *Reliability:* Because the gateway proxies all the appliances, any outages in communication due to the appliances roaming out of range are transparent to the host/server.
- *Longer battery life:* When the appliances are idle, the network software does not have to periodically send out keep-alive packets to keep the connection to the host/server open. The gateway does this.

Wireless Network Interface

Computers process information in digital form, with low direct current (DC) voltages representing data 1s and 0s. These signals are optimum for transmission within the computer, not for transporting data through wired or wireless media. A wireless network interface couples the digital signal from the end-user appliance to the wireless medium, which is air, to enable an efficient transfer of data between sender and receiver. This process includes the modulation and amplification of the digital signal to a form acceptable for propagation to the receiving location.

Note

Modulation is the process of translating the baseband digital signal used in the appliance to an analog form suitable for transmission through the air. This process is very similar to the common telephone modem, which converts a computer's digital data into an analog form within the 4 KHz limitation of the telephone circuit. The wireless modulator translates the digital signal to a frequency that propagates well through the atmosphere. Of course wireless networks employ modulation by using radio waves and infrared light.

The wireless network interface generally takes the shape of a wireless NIC or an external modem that facilitates the modulator and communications protocols. These components interface with the user appliance via a computer bus, such as ISA (Industry Standard Architecture) or PCMCIA (Personal Computer Memory Card International Association). The ISA bus comes standard in most desktop PCs. Many portable computers have PCMCIA slots that accept credit card-sized NICs. PCMCIA specifies three interface sizes: Type I (3.3 millimeters), Type II (5.0 millimeters), and Type III (10.5 millimeters). Some companies also produce wireless components that connect to the computer via the RS-232 serial port.

The interface between the user's appliance and NIC also includes a software driver that couples the client's application or NOS software to the card. The following driver standards are common:

- *NDIS (Network Driver Interface Specification)*: Driver used with Microsoft network operating systems
- *ODI (Open Datalink Interface)*: Driver used with Novell network operating systems
- *PDS (Packet Driver Specification)*: A generic DOS-based driver developed by FTP Software, Inc. for use with TCP/IP-based implementations

Note

Be sure to investigate the existence of suitable (NDIS, ODI, PACKET) drivers for the wireless NIC, and fully test its functionality with your chosen appliance before making large investments in wireless network hardware.

Radio cards traditionally come in a two-piece version configuration—that is, a PCMCIA card that inserts into the appliance and an external transceiver box. This setup is okay for some applications, such as forklift-mounted appliances; however, it is not ergonomic for most handheld appliances. Some vendors, especially with their newest radio cards, offer one-piece units having an integrated radio and transceiver assembly that all fits within the PCMCIA form factor.

Antenna

The antenna radiates the modulated signal through the air so that the destination can receive it. Antennas come in many shapes and sizes and have the following specific electrical characteristics:

- Propagation pattern
- Gain
- Transmit power
- Bandwidth

The *propagation pattern* of an antenna defines its coverage. A truly omnidirectional antenna transmits its power in all directions; whereas, a directional antenna concentrates most of its power in one direction. Figure 1.9 illustrates the differences.

A directional antenna has more *gain* (degree of amplification) than the omnidirectional type and is capable of propagating the modulated signal farther because it focuses the power in a single direction. The amount of gain depends on the directivity of the antenna. An omnidirectional antenna has a gain equal to one; that is, it doesn't focus the power in any particular direction. Omnidirectional antennas are best for indoor wireless networks because of relatively shorter range requirements and less susceptibility to outward interference.

Directional antennas will best satisfy needs for interconnecting buildings within metropolitan areas because of greater range and the desire to minimize interference with other systems.

The combination of *transmit power* and gain of an antenna defines the distance the signal will propagate. Long-distance transmissions require higher power and directive radiation patterns; whereas, shorter distance transmissions can get by with less power and gain. With wireless networks, the transmit power is relatively low, typically one watt or less.




FIGURE 1.9 An omnidirectional antenna broadcasts radio waves in all directions; whereas, a directional antenna focuses the power in a particular direction.

Note

Most spread spectrum radio vendors sell the following types of antennas:

- Snap-on antenna: Connects directly to the radio card and provides relatively low gain via an omnidirectional radio propagation pattern. This relatively small antenna is best for highly mobile applications when a larger antenna is impractical.
- Dipole antenna: Sits on a desk or table and connects to the radio card via a short antenna cable. This approach provides relatively low gain. This antenna is best for portable applications.
- High gain antenna: Attaches to a wall or antenna pole/tower and connects to the radio card or access point via a relatively long antenna cable. This approach provides relatively high gain and is best for access points and permanent stations.

Bandwidth is the effective part of the frequency spectrum that the signal propagates. The telephone system, for example, operates over a bandwidth roughly from 0 to 4 KHz. This is enough bandwidth to accommodate most of the frequency components within our voices. Radio wave systems have greater amounts of bandwidths located at much higher frequencies. Data rates and bandwidth are directly proportional: the higher the data rates, the more bandwidth you will need.

Note

If you're considering integrating a radio NIC into a particular PCMCIA-based appliance, such as a hand-held data collector, you may have to redesign the antenna mounting hardware to accommodate the construction of the appliance.

The Communications Channel

All information systems employ a communications channel along which information flows from source to destination. Ethernet networks may utilize twisted-pair or coaxial cable. Wireless networks use air as the medium. At the earth's surface, where most wireless networks operate, pure air contains gases, such as nitrogen and oxygen. This atmosphere provides an effective medium for the propagation of radio waves and infrared light.

Troubleshooting Tip

The communications channel offers unforeseen obstacles to wireless systems. Always perform a site survey to investigate the effects of physical structures and atmospheric conditions on the propagation of wireless signals before finalizing the design and purchase of a wireless system. (See "Identifying the Location of Access Points" in Chapter 8, "Implementing a Wireless LAN," for information on conducting a site survey.)

Rain, fog, and snow can increase the amount of water molecules in the air, however, and can cause significant *attenuation* to the propagation of modulated wireless signals. Smog clutters the air, adding attenuation to the communications channel as well. Attenuation is the decrease in the amplitude of the signal, and it limits the operating range of the system. The ways to combat attenuation are to either increase the transmit power of the wireless devices, which in most cases is limited by the FCC, or incorporate special amplifiers called *repeaters* that receive attenuated signals, revamp them, and transmit downline to the end station or next repeater.

Logical Architecture of a Wireless Network

A *logical architecture* defines the network's protocols, which ensures a well-managed and effective means of communication. PCs, servers, routers, and other active devices must conform to very strict rules to facilitate the proper coordination and transfer of information.

One popular standard logical architecture is the seven-layer Open System Interconnect (OSI) Reference Model, developed by the International Standards Organization (ISO). OSI specifies a complete set of network functions, grouped into layers. Figure 1.10 illustrates the OSI Reference Model.

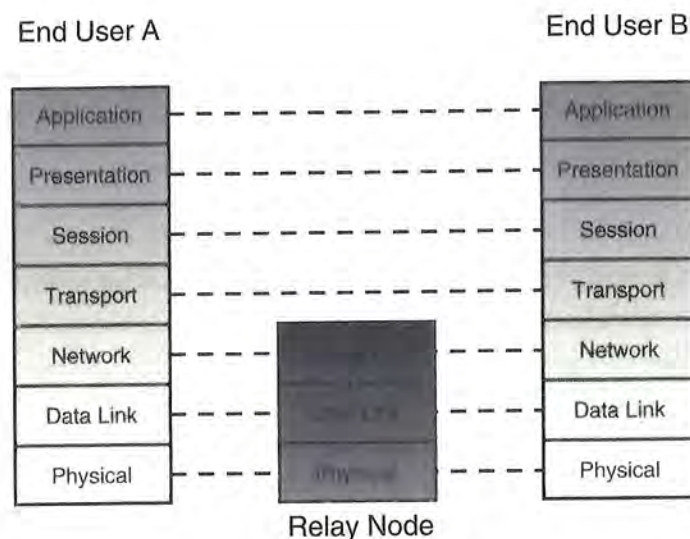


FIGURE 1.10 The Open System Interconnect Reference Model illustrates all levels of network functionality.

The OSI layers provide the following network functionality:

- *Layer 7—Application Layer:* Establishes communications with other users and provides such services as file transfer and email to the end users of the network.
- *Layer 6—Presentation Layer:* Negotiates data transfer syntax for the Application Layer and performs translations between different data types, if necessary.
- *Layer 5—Session Layer:* Establishes, manages, and terminates sessions between applications.
- *Layer 4—Transport Layer:* Provides mechanisms for the establishment, maintenance, and orderly termination of virtual circuits, while shielding the higher layers from the network implementation details. Such protocols as TCP operate at this layer.
- *Layer 3—Network Layer:* Provides the routing of packets through routers from source to destination. Such protocols as IP operate at this layer.
- *Layer 2—Data Link Layer:* Ensures synchronization and error control between two entities.
- *Layer 1—Physical Layer:* Provides the transmission of bits through a communication channel by defining electrical, mechanical, and procedural specifications.

Note

Each layer of OSI supports the layers above it.

Does a wireless network offer all OSI functions? No, not in a theoretical sense. As shown in Figure 1.11, wireless networks operate only within the bottom three layers. Only wireless wide area networks, however, perform Network Layer functions.

In addition to the wireless network functions, a complete network architecture needs to include such functions as end-to-end connection establishment and application services to make it useful. Chapter 3, “Overview of the IEEE 802.11 Standard,” provides details on the architecture of 802.11-compliant LANs which only covers the Network and Physical Layers of OSI. Chapter 6, “Wireless System Integration,” explains other components necessary to design and implement a complete system.

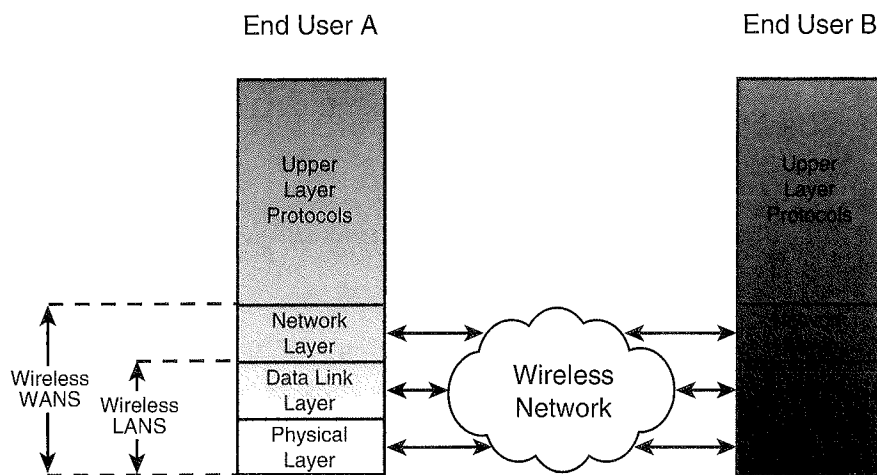


FIGURE 1.11 *Wireless LANs and MANs fulfill Data Link and Physical Layer functionality; whereas, wireless WANs also include functions at the Network Layer.*

The History of Wireless Networks

Network technologies and radio communications were brought together for the first time in 1971 at the University of Hawaii as a research project called ALOHANET. The ALOHANET system enabled computer sites at seven campuses spread out over four islands to communicate with the central computer on Oahu without using the existing unreliable and expensive phone lines. ALOHANET offered bidirectional communications, in a star topology, between the central computer and each of the remote stations. The remote stations had to communicate with one another via the centralized computer.

In the 1980s, amateur radio hobbyists, *hams*, kept radio networking alive within the United States and Canada by designing and building *terminal node controllers* (TNCs) to interface their computers through ham radio equipment (see Figure 1.12). TNCs

act much like a telephone modem, converting the computer's digital signal into one that a ham radio can modulate and send over the airwaves by using a packet-switching technique. In fact, the American Radio Relay League (ARRL) and the Canadian Radio Relay League (CRRL) have been sponsoring the Computer Networking Conference since the early 1980s to provide a forum for the development of wireless WANs. Thus, hams have been utilizing wireless networking for years, much earlier than the commercial market.

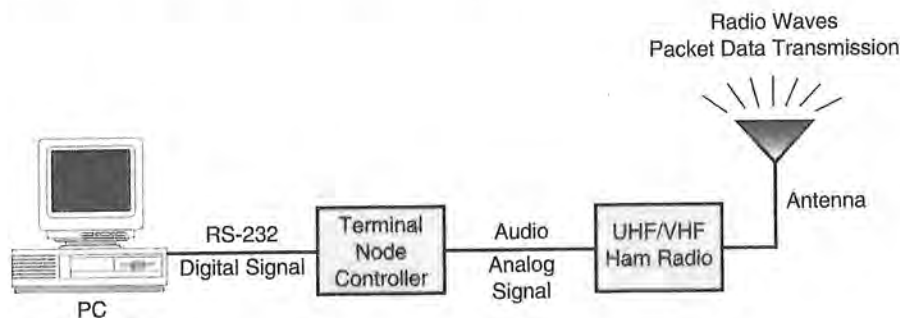


FIGURE 1.12 Terminal node controllers enable a PC to interface with a ham radio to form a packet radio network.

In 1985, the Federal Communications Commission (FCC) made the commercial development of radio-based LAN components possible by authorizing the public use of the Industrial, Scientific, and Medical (ISM) bands. This band of frequencies resides between 902 MHz and 5.85 GHz, just above the cellular phone operating frequencies. The ISM band is very attractive to wireless network vendors because it provides a part of the spectrum upon which to base their products, and end users do not have to obtain FCC licenses to operate the products. The ISM band allocation has had a dramatic effect on the wireless industry, prompting the development of wireless LAN components. Without a standard, however, vendors began developing proprietary radios and access points.

In the late 1980s, the Institute for Electrical and Electronic Engineers (IEEE) 802 Working Group, responsible for the development of LAN standards, such as ethernet and token ring, began development of standards for wireless LANs. Under the chairmanship of Vic Hayes, an engineer from NCR, the IEEE 802.11 Working Group developed the Wireless LAN Medium Access Control and Physical Layer specifications.

The IEEE Standards Board approved the standard on June 26, 1997, and the IEEE published the standard on November 18, 1997. The finalizing of this standard is prompting vendors to release 802.11-compliant radio cards and access points throughout 1998. Other vendors new to the wireless market are sure to develop and

release 802.11-compliant products based on the standard blueprint provided by the 802.11 standard.

Another widely accepted wireless network connection, however, has been wireless WAN services, which began surfacing in the early 1990s. Companies such as ARDIS and RAM Mobile Data were first in selling wireless connections between portable computers, corporate networks, and the Internet. Companies then began introducing Cellular Digital Packet Data (CDPD) services, which enable users to send and receive data packets via digital transmission services. These services enable employees to access email and other information services from their personal appliances without using the telephone system when meeting with customers, traveling in the car, or staying in a hotel.

The Future of Wireless Networks

Where is wireless networking going? What will the future bring? Predicting what the state of this technology and its products will be five years from now, or even a year from now, is impossible. The outlook for wireless networks, however, is very good. The maturation of standards should motivate vendors to produce new wireless products and drive the prices down to levels that are much easier to justify.

The presence of standards will motivate smaller companies to manufacture wireless components because they will not need to invest large sums of money in the research and development phases of the product. These investments already will have been made and embodied within the standards, which will be available to anyone interested in building wireless network components.

CHAPTER 2

Wireless Network Configurations

- **Wireless LANs**

It is important to understand the various types of wireless LANs to choose the best alternative technology and select the right components for use within a local area. You learn about the different configurations of a wireless LAN and how they operate.

- **Wireless point-to-point networks**

While providing network connectivity—mostly outdoors—wireless point-to-point networks offer additional challenges that are different from wireless LANs. Understanding how to maximize the use of wireless point-to-point network technologies is crucial to implementing their solutions.

- **Wireless WANs**

Wireless WANs can solve your wide area mobile network connectivity needs, but you need to use the technology that is going to provide the necessary coverage. You learn to differentiate the choices you have for wireless wide area networks.

Wireless LANs

Most wireless LANs operate over unlicensed frequencies at near-ethernet speeds (10 Mbps) using carrier sense protocols to share a radio wave or infrared light medium. The majority of these devices are capable of transmitting information up to 1,000 feet between computers within an open environment, and their costs per user range from \$150 to \$800. In addition, most wireless LAN products offer Simple Network Management Protocol (SNMP) to support network management through the use of SNMP-based management platforms and applications. Figure 2.1 illustrates the concept of a wireless local area network interfacing with a wired network.




FIGURE 2.1 A wireless local area network provides connectivity over the airwaves within a local area, such as a building.

The components of a wireless LAN consist of a wireless NIC and a wireless local bridge, which is often referred to as an *access point*. The wireless NIC interfaces the appliance with the wireless network, and the access point interfaces the wireless network with a wired network. Most wireless NICs interface appliances to the wireless network by implementing a carrier sense access protocol and modulating the data signal with a spreading sequence.

The following sections describe three approaches to wireless networking within a local environment. These methods include the following:

- Radio waves
- Infrared light
- Carrier currents

Radio-Based Wireless LANs

The most widely sold wireless LAN products use radio waves as a medium between computers and peripherals. An advantage of radio waves over other forms of wireless

connectivity is that they can interconnect users without line of sight and propagate through walls and other obstructions with fairly little attenuation, depending on the type of wall construction. Although several walls might separate the user from the server or wireless bridge, users can maintain connections to the network. This supports true mobility. With radio-LAN products, a user with a portable computer can move freely through the facility while accessing data from a server or running an application.

A disadvantage of using radio waves, however, is that an organization must manage them along with other electromagnetic propagation. Medical equipment and industrial components may utilize the same radio frequencies as wireless LANs, which could cause interference. An organization must determine whether potential interference is present before installing a radio-based LAN. Because radio waves penetrate walls, security might also be a problem. Unauthorized people from outside the controlled areas could receive sensitive information; however, vendors often scramble the data signal to protect the information from being understood by inappropriate people.

This section discusses the following topics that explain the operation and configuration of radio-based wireless LANs:

- Medium access control
- Spread spectrum modulation
- Narrowband modulation
- Wireless local bridges
- Single-cell wireless LANs
- Multiple-cell wireless LANs

Medium Access Control

Medium access control, which is a Data Link Layer function in a radio-based wireless LAN, enables multiple appliances to share a common transmission medium via a carrier sense protocol similar to ethernet. This protocol enables a group of wireless computers to share the same frequency and space.

As an analogy, consider a room of people engaged in a single conversation in which each person can hear if someone speaks. This represents a fully connected bus topology (where everyone communicates using the same frequency and space) that ethernet and wireless networks, especially wireless LANs, utilize.

To avoid having two people speak at the same time, you should wait until the other person has finished talking. Also, no one should speak unless the room is silent.

CHAPTER 3

Overview of the IEEE 802.11 Standard

- **The importance of standards**
This chapter begins with an introduction to the types of LAN standards and the primary organization that makes the standards: the Institute for Electrical and Electronic Engineers (IEEE). You learn the important benefits of using the IEEE 802.11 wireless LAN standard.
- **IEEE 802 LAN standards family**
It is important to know how the IEEE 802.11 standard fits into other LAN protocols to ensure proper interoperability. An overview of the 802 series of LAN standards describes the operation of the 802.2 Logical Link Control that directly interfaces with 802.11.
- **Introduction to the IEEE 802.11 standard**
An explanation of the scope and goals of the 802.11 standard provides an understanding of the basic functionality of 802.11. Learn the peculiar wireless network issues that were addressed when developing the standard.
- **IEEE 802.11 topology**
An overview of the physical structure of 802.11-compliant LANs provides an understanding of 802.11 topology. Understand how basic physical 802.11 elements, such as Basic Service Sets (single-cell wireless LANs) and access points, form integrated, multiple-cell wireless LANs that support a variety of mobility types.
- **IEEE 802.11 logical architecture**
Coverage of the main elements of the 802.11 protocol stack provides an overview of how the 802.11 protocol works. Learn the main functionality of each of the following 802.11 protocol layers: MAC Layer and individual PHY (Physical) Layers (frequency hopping, direct sequence, and infrared).

- **IEEE 802.11 services**
802.11-compliant LANs function based on a set of services that relate to stations and distribution systems. Discover how these services offer security equivalent to wired LANs.
- **Implications of the IEEE 802.11 standard**
Although the long-awaited 802.11 standard offers several benefits over using proprietary-based wireless LANs, the 802.11 standard still has shortcomings that implementors should be aware of. Learn some of the 802.11 implications, such as relatively low data rates and lack of roaming.
- **IEEE 802.11 standard compliance**
The compliance with 802.11 depends on those having the need for wireless networks. Become aware of how vendors are complying with 802.11, what end users need to do to be compliant, and how different regions of the world comply with 802.11 radio frequencies.
- **IEEE 802.11 Working Group operations**
Involvement in IEEE 802.11 standards development is open to anyone with a desire to participate, but you need to understand the membership requirements and types of 802.11 members.
- **Future of the IEEE 802.11 standard**
When making decisions about wireless LANs, be sure to include what the future holds for the 802.11 standard. Discover the projects IEEE 802.11 members are working on to increase the performance of 802.11-compliant wireless LANs.

The Importance of Standards

Vendors and some end users initially expected markets to dive headfirst into implementing wireless networks. Markets did not respond as predicted, and flat sales growth of wireless networking components prevailed through most of the 1990s. Relatively low data rates, high prices, and especially the lack of standards kept many end users from purchasing the wire-free forms of media.

For those having applications suitable for lower data rates and enough cost savings to warrant purchasing wireless connections, the only choice before 1998 was to install proprietary hardware to satisfy requirements. As a result, many organizations today have proprietary wireless networks for which you have to replace both hardware and software to be compliant with the IEEE 802.11 standard. The lack of standards has been a significant problem with wireless networking, but the first official version of the standard is now available. In response to lacking standards, the Institute for Electrical and Electronic Engineers (IEEE) developed the first internationally recognized wireless LAN standard: IEEE 802.11.

Types of Standards

There are two main types of standards: official and public. An *official standard* is published and known to the public, but it is controlled by an official standards organization, such as IEEE. Government or industry consortiums normally sponsor official standards groups. Official standards organizations generally ensure coordination at both the international and domestic level.

A *public standard* is similar to an official standard, except it is controlled by a private organization, such as the Wireless LAN Interoperability Forum. Public standards, often called *de facto standards*, are common practices that have not been produced or accepted by an official standards organization. These standards, such as TCP/IP, are the result of widespread proliferation. In some cases, public standards that proliferate, such as the original ethernet, eventually pass through standards organizations and become official standards.

Companies should strive to adopt standards and recommended products within their organizations for all aspects of information systems. What type of standards should you use? For most cases, focus on the use of an official standard if one is available and proliferating. This will help ensure widespread acceptance and longevity of your wireless network implementation. If no official standard is suitable, a public standard would be a good choice. In fact, public standards can often respond faster to changes in market needs because they usually have less organizational overhead for making changes. Be sure to avoid nonstandard or proprietary system components, unless there are no suitable standards available.

Case Study 3.1: 802.11 Versus Proprietary Standards

A large retail chain based in Sacramento, California, had requirements to implement a wireless network to provide mobility within their 10 warehouses located all over the United States. The application calls for clerks within the warehouse to utilize new handheld wireless data collectors that perform inventory-management functions.

The company, already having one vendor's data collection devices (we'll call these brand X), decides to use that vendor's

brand Y proprietary wireless data collectors and their proprietary wireless network (the vendor doesn't offer an 802.11-compliant solution). This decision eliminates the need to work with additional vendors for the new handheld devices and the wireless network.

A year passes since the installation, and enhancement requirements begin to pour in for additional mobile appliances that are not available from the brand X

continues

continued

vendor. This forces the company to consider the purchase of new brand Z appliances from a different vendor. The problem, however, is that the brand Z appliances, which are 802.11-compliant, don't interoperate with the installed proprietary brand Y wireless network. Because of the cost associated with replacing their network with one that is 802.11 compliant (the brand Y wireless network has no upgrade path to 802.11), the company can't cost effectively implement the new enhancement.

The company could have eliminated the problem of not being able to implement the new enhancement if it would have implemented the initial system with 802.11-compliant network components, because most vendors offer products that are compatible with 802.11, but not all the proprietary networks. The result would have been the ability to consider multiple vendors for a wider selection of appliances.

Institute for Electrical and Electronic Engineers (IEEE)

The IEEE is a nonprofit professional organization founded by a handful of engineers in 1884 for the purpose of consolidating ideas dealing with electro-technology. In the last 100 plus years, IEEE has maintained a steady growth. Today, the IEEE, which is based in the United States, has over 320,000 members located in 150 countries. The IEEE consists of 35 individual societies, including the Communications Society, Computer Society, and Antennas and Propagation Society.

The IEEE plays a significant role in publishing technical works, sponsoring conferences and seminars, accreditation, and standards development. The IEEE has published nearly 700 active standards publications, half of which relate to power engineering and most others deal with computers. The IEEE standards development process consists of 30,000 volunteers (who are mostly IEEE members) and a Standards Board of 32 people. In terms of LANs, IEEE has produced some very popular and widely used standards. The majority of LANs in the world utilize network interface cards based on the IEEE 802.3 (ethernet) and IEEE 802.5 (token ring) standards, for example.

Before someone can develop an IEEE standard, he must submit a Project Authorization Request (PAR) to the IEEE Standards Board. If the board approves the PAR, IEEE establishes a standards working group to develop the standard. Members of the working groups serve voluntarily and without compensation, and they are not necessarily members of the institute. The working group begins by writing a draft standard, and then solicits the draft to a balloting group of selected IEEE members for review and approval. The ballot group consists of the standard's developers, potential users, and other people having general interest.

Before publication, the IEEE Standards Board performs a review of the Final Draft Standard, and then considers approval of the standard. The resulting standard represents a consensus of broad expertise from within IEEE and other related organizations. All IEEE standards are subjected to review at least once every five years for revision or reaffirmation.

Note

In May 1991, a group of people, led by Victor Hayes, submitted a Project Authorization Request (PAR) to IEEE to initiate the 802.11 Working Group. Victor became Chairman of the working group and led the standards effort to its completion in June 1997.

Benefits of the 802.11 Standard

The benefits of utilizing standards, such as those published by IEEE, are great. The following sections explain the benefits of complying with standards, especially IEEE 802.11.

Appliance Interoperability

Compliance with the IEEE 802.11 standard makes interoperability between multiple-vendor appliances and the chosen wireless network type possible. This means you can purchase an 802.11-compliant PalmPilot from Symbol and Pathfinder Ultra handheld scanner/printer from Monarch Marking Systems, and they will both interoperate within an equivalent 802.11 wireless network, assuming 802.11 configuration parameters are set equally in both devices. Standard compliance increases price competition and enables companies to develop wireless LAN components with lower research and development budgets. This enables a greater number of smaller companies to develop wireless components. As a result, the sales of wireless LAN components should boom over the next few years as the finalization of the IEEE 802.11 standard sinks in.

As shown in Figure 3.1, appliance interoperability avoids the dependence on a single vendor for appliances. Without a standard, for example, a company having a non-standard proprietary Symbol network would be dependent on purchasing only appliances that operate on a Symbol network. This would exclude appliances such as ones from Telxon that only operate on proprietary Aironet networks. With an 802.11-compliant wireless network, you can utilize any equivalent 802.11-compliant appliance. Because most vendors, including Symbol and Telxon, have migrated their products to 802.11, you have a much greater selection of appliances for 802.11 standard networks.




FIGURE 3.1 *Appliance interoperability ensures that multiple-vendor hardware works within equivalent wireless networks.*

Fast Product Development

The 802.11 standard is a well-tested blueprint that developers can use to implement wireless devices. The use of standards decreases the learning curve required to understand specific technologies because the standard-forming group has already invested the time to smooth out any wrinkles in the implementation of the applicable technology. This leads to the development of products in much less time.

Stable Future Migration

Compliance with standards helps protect investments and avoids legacy systems that must be completely replaced in the future as those proprietary products become obsolete. The evolution of wireless LANs should occur in a similar fashion as 802.3, ethernet. Initially, ethernet began as a 10 Mbps standard using coaxial cable media. The IEEE 802.3 Working Group enhanced the standard over the years by adding twisted-pair, optical-fiber cabling, and 100 and 1000 Mbps data rates.

Just as IEEE 802.3 did, the 802.11 Working Group recognizes the investments organizations make in network infrastructure and the importance in providing migration paths that maximize the installed base of hardware. As a result, 802.11 will certainly ensure stable migration from existing wireless LANs as higher performance wireless networking technologies become available.

Price Reductions

High costs have always plagued the wireless LAN industry; however, prices should drop significantly as more vendors and end users comply with 802.11. One of the reasons for lower prices is that vendors will no longer need to develop and support lower-quantity proprietary subcomponents, cutting design, manufacturing, and support costs. Ethernet went through a similar lowering of prices as more and more companies began complying with the 802.3 standard.

Avoiding Silos

Over the past couple of decades, MIS organizations have had a difficult time maintaining control of network implementations. The introduction of PCs, LANs, and visual-based development tools has made it much easier for non-MIS organizations, such as finance and manufacturing departments, to deploy their own applications. One part of the company, for example, may purchase a wireless network from one vendor, and then another part of the company may buy a different wireless network. As a result, *silos*—noninteroperable systems—appear within the company, making it very difficult for MIS personnel to plan and support compatible systems. Some people refer to these silos as *stovepipes*.

Acquisitions bring dissimilar systems together as well. One company having a proprietary system may purchase another having a different proprietary system, resulting in noninteroperability. Figure 3.2 illustrates the features of standards that minimize the occurrence of silos.

Case Study 3.2:

Problems with Mixed Standards

A company located in Barcelona, Spain specializes in the resale of women's clothes. This company, having a MIS group without much control over the implementation of distributed networks in major parts of the company, has projects underway to implement wireless networks for an inventory application and a price-marking application.

Non-MIS project managers located in different parts of the company lead these projects. They have little desire to coordinate their projects with MIS because of past difficulties. As a result, both pro-

ject managers end up implementing noncompatible proprietary wireless networks to satisfy their networking requirements.

The project managers install both systems: one that covers the sales floorspace of their 300 stores (for price marking) and one that encompasses 10 warehouses (for doing inventory functions). Although the systems are noncompatible, all is fine for the users operating the autonomous systems.

The issues with this system architecture, however, are the difficulty in providing

continues

continued

operational support and inflexibility. The company must maintain purchasing and warranty contracts with two different wireless network vendors, service personnel need to acquire and maintain an understanding in the operation of two networks, and the company cannot share appliances and wireless network components between the warehouses and the stores.

As a result, the silos in this case make the networks more expensive to support and limit their flexibility in meeting future needs. The implementation of standard 802.11-compliant networks would have avoided these problems.




FIGURE 3.2 Compliance with the IEEE 802.11 standard can minimize the implementation of silos.

IEEE 802 LAN Standards Family

The IEEE 802 Local and Metropolitan Area Network Standards Committee is a major working group chartered by IEEE to create, maintain, and encourage the use of IEEE and equivalent IEC/ISO standards. IEEE formed the committee in February 1980, and has met at least three times per year as a plenary body since then. IEEE 802 produces the series of standards known as IEEE 802.x, and the JTC 1 series of equivalent standards are known as ISO 8802-nnn.

IEEE 802 includes a family of standards, as depicted in Figure 3.3. The MAC and Physical Layers of the 802 standard were organized into a separate set of standards from the LLC because of the interdependence between medium access control, medium, and topology.

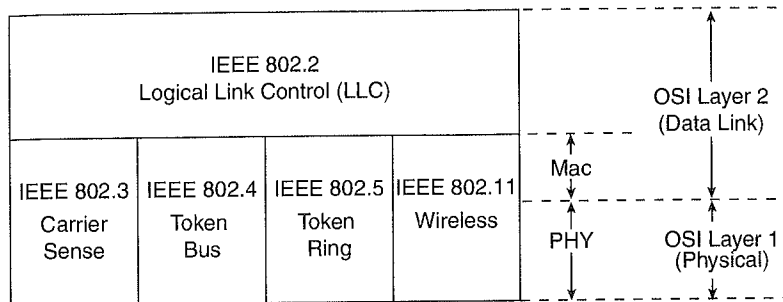


FIGURE 3.3 The IEEE 802 family of standards falls within the scope of layers 1 and 2 of the OSI Reference Model. The LLC protocol specifies the mechanisms for addressing stations across the medium and for controlling the exchange of data between two stations; whereas, the MAC and PHY Layers provide medium access and transmission functions.

The IEEE 802 family of standards includes the following:

- *IEEE 802.1: Glossary, Network Management, and Internetworking:* These documents, as well as IEEE 802 Overview and Architecture, form the scope of work for the 802 standards.
- *IEEE 802.2: Logical Link Control (LLC):* This standard defines Layer 2 synchronization and error control for all types of 802 LANs, including 802.11. Refer to the next section, “IEEE 802.2 LLC Overview,” for more detail on the features and operation of the LLC.
- *IEEE 802.3: CSMA/CD Access Method and Physical Layer Specifications:* This defines the widely accepted 10, 100, and 1000 Mbps ethernet asynchronous protocol for use over twisted-pair wiring, coaxial cable, and optical fiber.
- *IEEE 802.4: Token-Passing Bus Access Method and Physical Layer Specifications:* This offers a token-passing protocol over a bus topology that can be embedded in other systems.
- *IEEE 802.5: Token-Passing Ring Access Method and Physical Layer Specifications:* This defines a 4 and 16 Mbps synchronous protocol that uses a token for access control over a ring topology.
- *IEEE 802.10: Security and Privacy Access Method and Physical Layer Specifications:* Provides security provisions for both wired and wireless LANs.
- *IEEE 802.11: Wireless Access Method and Physical Layer Specification:* Encompasses a variety of physical media, including frequency hopping spread spectrum, direct sequence spread spectrum, and infrared light for data rates up to 2 Mbps.

IEEE 802.2 LLC Overview

The LLC is the highest layer of the IEEE 802 Reference Model and provides similar functions of the traditional Data Link Control protocol: HDLC (High-Level Data

Link Control). The ANSI/IEEE Standard 802.2 specifies the LLC. The purpose of the LLC is to exchange data between end users across a LAN using a 802-based MAC controlled link. The LLC provides addressing and data link control, and it is independent of the topology, transmission medium, and medium access control technique chosen.

Higher layers, such as TCP/IP, pass user data down to the LLC expecting error-free transmission across the network. The LLC in turn appends a control header, creating an LLC protocol data unit (PDU). The LLC utilizes the control information in the operation of the LLC protocol (see Figure 3.4). Before transmission, the LLC PDU is handed down through the MAC service access point (SAP) to the MAC Layer, which appends control information at the beginning and end of the packet, forming a MAC frame. The control information in the frame is needed for the operation of the MAC protocol.

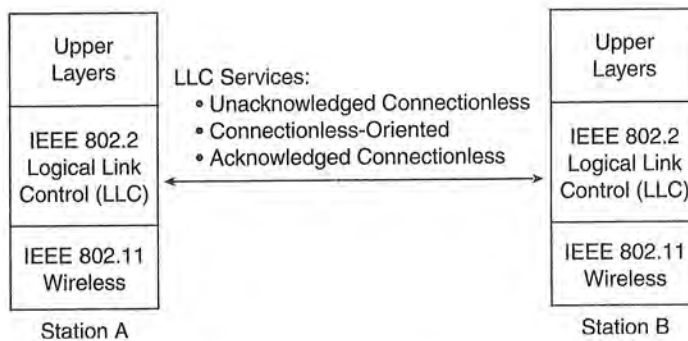


FIGURE 3.4 The LLC provides end-to-end link control over an 802.11-based wireless LAN.

IEEE 802.2 LLC Services

The LLC provides the following three services for a Network Layer protocol:

- Unacknowledged connectionless service
- Connection-oriented service
- Acknowledged connectionless service

These services apply to the communication between peer LLC Layers—that is, one located on the source station and one located on the destination station. Typically, vendors will provide these services as options that the customer can select when purchasing the equipment.

All three LLC protocols employ the same PDU format that consists of four fields (see Figure 3.5). The Destination Service Access Point (DSAP) and Source Service Access Point (SSAP) fields each contains 7-bit addresses, which specify the destination and

source stations of the peer LLCs. One bit of the DSAP indicates whether the PDU is intended for an individual or group station(s). One bit of the SSAP indicates whether it is a command or response PDU. The format of the LLC Control field is identical to that of HDLC, using extended (7-bit) sequence numbers. The Data field contains the information from higher-layer protocols that the LLC is transporting to the destination.

8 Bits	8 Bits	8 Bits	Variable
Destination SAP	Service SAP	Control	Data

FIGURE 3.5 The LLC PDU consists of data fields that provide the LLC functionality.

The Control field has bits that indicate whether the frame is one of the following types:

- *Information:* Used to carry user data
- *Supervisory:* Used for flow control and error control
- *Unnumbered:* Various protocol control PDUs

Unacknowledged Connectionless Service

The *unacknowledged connectionless service* is a datagram-style service that does not involve any error-control or flow-control mechanisms. This service does not involve the establishment of a Data Link Layer connection (that is, a connection between peer LLCs). This service supports individual, multicast, and broadcast addressing. This service just sends and receives LLC PDUs, with no acknowledgment of delivery. Because the delivery of data is not guaranteed, a higher layer, such as TCP, must deal with reliability issues.

The unacknowledged connectionless service offers advantages in the following situations:

- If higher layers of the protocol stack provide the necessary reliability and flow-control mechanisms, it would be inefficient to duplicate them in the LLC. In this case, the unacknowledged connectionless service would be appropriate. TCP and the ISO transport protocol, for example, already provide the mechanisms necessary for reliable delivery.
- It is not always necessary to provide feedback pertaining to successful delivery of information. The overhead of connection establishment and maintenance can be inefficient—as an example, for applications involving the periodic sampling of data sources, such as monitoring sensors. The unacknowledged connectionless service would best satisfy these requirements.

**Case Study 3.3:
Using Unacknowledged
Connectionless Service to
Minimize Overhead**

The executive office building of a high-rent advertising agency in Southern California has 20 sensors to monitor temperatures throughout its building as an input to the heating and air conditioning system. These sensors send short information packets every minute to an application on a centralized server that updates a temperature table in a database. The heating and air conditioning system uses this information to control the temperature in different parts of the building.

For this application, the server does not need to acknowledge the reception of

every sensor transmission because the information updates are not critical. The system can maintain a comfortable temperature throughout the building even if the system misses temperature updates from time to time.

Additionally, it is not feasible to require the sensors to establish connections with the server to send the short information packets. As a result, designers of the system chose to use the LLC unacknowledged connectionless service to minimize overhead on the network, making the limited wireless network bandwidth available to other applications.

Connection-Oriented Service

The *connection-oriented service* establishes a logical connection that provides flow control and error control between two stations needing to exchange data. This service does involve the establishment of a connection between peer LLCs by performing connection establishment, data transfer, and connection termination functions. The service can only connect two stations; therefore, it does not support multicast or broadcast modes. The connection-oriented service offers advantages mainly if higher layers of the protocol stack do not provide the necessary reliability and flow-control mechanisms, which is generally the case with terminal controllers.

Flow control is a protocol feature that ensures a transmitting station does not overwhelm a receiving station with data. With flow control, each station allocates a finite amount of memory and buffer resources to store sent and received PDUs.

Networks, especially wireless networks, suffer from induced noise in the links between network stations that can cause transmission errors. If the noise is high enough in amplitude, it causes errors in digital transmission in the form of altered bits. This will lead to inaccuracy of the transmitted data, and the receiving network device may misinterpret the meaning of the information.

The noise that causes most problems with networks is usually Gaussian and impulse noise. Theoretically, the amplitude of Gaussian noise is uniform across the frequency spectrum, and it normally triggers random single-bit independent errors.

Impulse noise, the most disastrous, is characterized by long quiet intervals of time followed by high-amplitude bursts. This noise results from lightning and switching transients. Impulse noise is responsible for most errors in digital communication systems and generally provokes errors to occur in bursts.

To guard against transmission errors, the connection-oriented and acknowledged-connectionless LLCs use error-control mechanisms that detect and correct errors that occur in the transmission of PDUs. The LLC ARQ mechanism recognizes the possibility of the following two types of errors:

- *Lost PDU*: A PDU fails to arrive at the other end or is damaged beyond recognition.
- *Damaged PDU*: A PDU has arrived, but some bits are altered.

When a frame arrives at a receiving station, the station checks whether there are any errors present by using a *Cyclic Redundancy Check* (CRC) error detection algorithm. In general, the receiving station will send back a positive or negative acknowledgment depending on the outcome of the error detection process. In case the acknowledgment is lost en route to the sending station, the sending station will retransmit the frame after a certain period of time. This process is often referred to as *Automatic Repeat-Request* (ARQ).

Overall, ARQ is best for the correction of burst errors because this type of impairment occurs in a small percentage of frames, thus not invoking many retransmissions. Because of the feedback inherent in ARQ protocols, the transmission links must accommodate half-duplex or full-duplex transmissions. If only simplex links are available due to feasibility, it is impossible to use the ARQ technique because the receiver would not be able to notify the transmitter of bad data frames.

Note

In cases for which single bit errors predominate or when only a simplex link is available, forward error correction (FEC) can provide error correction. FEC algorithms provide enough redundancy in data transmissions to enable the receiving station to correct errors without needing the sending station to retransmit the data.

FEC is effective for correcting single-bit errors, but it requires a great deal of overhead in the transmissions to protect against multiple errors, such as burst errors. The IEEE LLC, however, specifies only the use of ARQ-based protocols for controlling errors.

The following are two approaches for retransmitting unsatisfactory blocks of data using ARQ.

Continuous ARQ

With continuous ARQ, often called a *sliding window protocol*, the sending station transmits frames continuously until the receiving station detects an error. The

sending station is usually capable of transmitting a specific number of frames and maintains a table indicating which frames have been sent.

The system implementor can set the number of frames sent before stopping via configuration parameters of the network device. If a receiver detects a bad frame, it will send a negative acknowledgment back to the sending station requesting that the bad frame be sent over again. When the transmitting station gets the signal to retransmit the frame, several subsequent frames may have already been sent (due to propagation delays between the sender and receiver); therefore, the transmitter must “go back” and retransmit the erred data frame.

There are a couple ways the transmitting station can send frames again using continuous ARQ. One method is for the source to retrieve the erred frame from the transmit buffer and send the bad frame and all frames following it. This is called the *go-back-n technique*. A problem, however, is when n (the number of frames the transmitter sent after the erred frame plus one) becomes large, the method becomes inefficient. This is because the retransmission of just one frame means that a large number of possibly “good” frames will also be resent, thus decreasing throughput.

The go-back- n technique is useful in applications for which receiver buffer space is limited because all that is needed is a receiver window size of one (assuming frames are to be delivered in order). When the receive node rejects an erred frame (sends a negative acknowledgment), it does not need to buffer any subsequent frames for possible reordering while it is waiting for the retransmission because all subsequent frames will also be sent.

An alternative to the continuous go-back- n technique is a method that selectively retransmits only the erred frame, and then resumes normal transmission at the point just before getting the notification of a bad data frame. This approach is called *selective repeat*. It is obviously better than continuous go-back- n in terms of throughput because only the erred frame needs retransmission. With this technique, however, the receiver must be capable of storing a number of frames if they are to be processed in order. The receiver needs to buffer data that have been received after an erred frame was requested for retransmission because only the damaged frame will be sent again.

Stop-and-Wait ARQ

With stop-and-wait ARQ, the sending station transmits a frame and then stops and waits for some type of acknowledgment from the receiver on whether a particular frame was acceptable or not. If the receiving station sends a negative acknowledgment, the frame will be sent again. The transmitter will send the next frame only after it receives a positive acknowledgment from the receiver.

An advantage of stop-and-wait ARQ is that it does not require much buffer space at the sending or receiving station. The sending station needs to store only the current transmitted frame. However, stop-and-wait ARQ becomes inefficient as the propagation delay between source and destination becomes large. For example, data sent on satellite links normally experience a round-trip delay of several hundred milliseconds; therefore, long block lengths are necessary to maintain a reasonably effective data rate. The trouble is that with longer frames, the probability of an error occurring in a particular block is greater. Therefore, retransmission will occur often, and the resulting throughput will be lower.

Case Study 3.4:
Using Automatic Repeat-Request (ARQ) to Reduce Errors

A mobile home manufacturer in Florida uses robots on the assembly line to perform welding. Designers of the robot control system had to decide whether to use ARQ or FEC for controlling transmission errors between the server and the robots. The company experiences a great deal of impulse noise from arc welders and other heavy machinery.

In the midst of this somewhat hostile environment, the robots require error-free information updates to ensure that they function correctly. Designers of the system quickly ruled out the use of FEC because of the likely presence of burst errors due to impulse noise. ARQ, with its capability to detect and correct frames having a lot of bit errors, was obviously the better choice.

Acknowledged Connectionless Service

As with the unacknowledged connectionless service, the *acknowledged connectionless service* does not involve the establishment of a logical connection with the distant station. But the receiving stations with the acknowledged version do confirm successful delivery of datagrams. Flow and error control is handled through use of the stop-and-wait ARQ method.

The acknowledged connectionless service is useful in several applications. The connection-oriented service must maintain a table for each active connection for tracking the status of the connection. If the application calls for guaranteed delivery, but there are a large number of destinations needing to receive the data, the connection-oriented service may be impractical because of the large number of tables required. Examples that fit this scenario include process control and automated factory environments that require a central site to communicate with a large number of processors and programmable controllers. In addition, the handling of important and time-critical alarm or emergency control signals in a factory would also fit this

case. In all these examples, the sending stations need an acknowledgment to ensure successful delivery of the data; however, the urgency of transmission cannot wait for a connection establishment.

Note

A company having a requirement to send information to multiple devices needing positive acknowledgment of the data transfer can make use of the acknowledged connectionless LLC service. A marina may find it beneficial to control the power to different parts of the boat dock via a wireless network, for example. Of course, the expense of a wireless network may not be justifiable for this application alone.

Other applications, such as supporting data transfers back and forth to the cash register at the gas pump and the use of data-collection equipment for inventorying rental equipment, can share the wireless network to make a more positive business case. For shutting off the power on the boat dock, the application would need to send a message to the multiple power controllers, and then expect an acknowledgment to ensure the controller receives the notification and that the power is shut off. For this case, the connectionless transfer, versus connection-oriented, makes most sense because it would not be feasible to make connections to the controllers to support such a short message.

LLC/MAC Layer Service Primitives

Layers within the 802 architecture communicate with each other via service primitives having the following forms:

- *Request:* A layer uses this type of primitive to request that another layer perform a specific service.
- *Confirm:* A layer uses this type of primitive to convey the results of a previous service request primitive.
- *Indication:* A layer uses this type of primitive to indicate to another layer that a significant event has occurred. This primitive could result from a service request or from some internally generated event.
- *Response:* A layer uses this type of primitive to complete a procedure initiated by an indication primitive.

These primitives are an abstract way of defining the protocol, and they *do not* imply a specific physical implementation method. Each layer within the 802 model uses specific primitives. The LLC communicates with its associated MAC Layer through the following specific set of service primitives:

- **MA-UNITDATA.request:** The LLC sends this primitive to the MAC Layer to request the transfer of a data frame from a local LLC entity to a specific peer LLC entity or group of peer entities on different stations. The data frame could be an information frame containing data from a higher layer or a control frame (for example, a supervisory or unnumbered frame) that the LLC generates internally to communicate with its peer LLC.
- **MA-UNITDATA.indication:** The MAC Layer sends this primitive to the LLC to transfer a data frame from the MAC Layer to the LLC. This occurs only if the

MAC has found that a frame it receives from the Physical Layer is valid, has no errors, and that the destination address indicates the correct MAC address of the station.

- *MA-UNITDATA-STATUS.indication*: The MAC Layer sends this primitive to the LLC Layer to provide status information about the service provided for a previous MA-UNITDATA.request primitive.

Note

The current ANSI/IEEE 802.2 standard (dated May 7, 1998) states that the 802.2 Working Group is developing a single-service specification of primitives that is common to all MAC Layers. IEEE will refer to this change in the 802.2 standard, not the individual MAC Layer standards (for example, 802.3, 802.5, 802.11).

Introduction to the IEEE 802.11 Standard

The initial 802.11 PAR states, "...the scope of the proposed [wireless LAN] standard is to develop a specification for wireless connectivity for fixed, portable, and moving stations within a local area." The PAR further says that the "purpose of the standard is to provide wireless connectivity to automatic machinery and equipment or stations that require rapid deployment, which may be portable, handheld, or which may be mounted on moving vehicles within a local area."

The resulting standard, which is officially called *IEEE Standard for Wireless LAN Medium Access (MAC) and Physical Layer (PHY) Specifications*, defines over-the-air protocols necessary to support networking in a local area. As with other IEEE 802-based standards (for example, 802.3 and 802.5), the primary service of the 802.11 standard is to deliver MSDUs (MAC Service Data Units) between peer LLCs. Typically, a radio card and access point provide functions of the 802.11 standard.

Note

To order a copy of the IEEE 802.11 standard, contact the IEEE 802 Document Order Service at 800-678-4333. You can also order the standard via IEEE's Web site at www.ieee.org.

The 802.11 standard provides MAC and PHY functionality for wireless connectivity of fixed, portable, and moving stations moving at pedestrian and vehicular speeds within a local area. Specific features of the 802.11 standard include the following:

- Support of asynchronous and time-bounded delivery service
- Continuity of service within extended areas via a distribution system, such as ethernet
- Accommodation of transmission rates of 1 and 2 Mbps
- Support of most market applications

- Multicast (including broadcast) services
- Network management services
- Registration and authentication services

Target environments for use of the standard include the following:

- Inside buildings, such as offices, banks, shops, malls, hospitals, manufacturing plants, and residences
- Outdoor areas, such as parking lots, campuses, building complexes, and outdoor plants

The 802.11 standard takes into account the following significant differences between wireless and wired LANs:

- *Power management:* Because most wireless LAN NICs are available in PCMCIA Type II format, obviously you can outfit portable and mobile handheld computing equipment with wireless LAN connectivity. The problem, however, is these devices must rely on batteries to power the electronics within them. The addition of a wireless LAN NIC to a portable computer can quickly drain batteries.

The 802.11 Working Group struggled with finding solutions to conserve battery power; however, they found techniques enabling wireless NICs to switch to lower-power standby modes periodically when not transmitting, reducing the drain on the battery. The MAC Layer implements power-management functions by putting the radio to sleep (that is, lowering the power drain) when no transmission activity occurs for some specific or user-definable time period. The problem, however, is that a sleeping station can miss critical data transmissions. 802.11 solves this problem by incorporating buffers to queue messages. The standard calls for sleeping stations to awaken periodically and retrieve any applicable messages.

- *Bandwidth:* The ISM spread spectrum bands do not offer a great deal of bandwidth, keeping data rates lower than desired for some applications. The 802.11 Working Group, however, dealt with methods to compress data, making the best use of available bandwidth. Efforts are also underway to increase the data rate of 802.11 to accommodate the growing need for exchanging larger and larger files (see the section titled “Future of the IEEE 802.11 Standard” at the end of this chapter).
- *Security:* As mentioned in Chapter 1, “Introduction to Wireless Networks,” in the “Network Security” section, wireless LANs transmit signals over much larger areas than that of wired media, such as twisted-pair, coaxial, and optical fiber cable. In terms of privacy, therefore, wireless LANs have a much larger area to protect. To employ security, the 802.11 Working Group coordinated their work with the IEEE 802.10 Standards Committee responsible for developing security mechanisms for all 802 series LANs.

- *Addressing:* The topology of a wireless network is dynamic; therefore, the destination address does not always correspond to the destination's location. This raises a problem when routing packets through the network to the intended destination. Therefore, you may need to utilize a TCP/IP-based protocol, such as MobileIP, to accommodate mobile stations. Chapter 6, "Wireless System Integration," provides details on the MobileIP protocol.

To ensure interoperability with existing standards, the 802.11 Working Group developed the standard to be compatible with other existing 802 standards, such as the following:

- *IEEE 802:* Functional Requirements
- *IEEE 802.2:* MAC Service Definition
- *IEEE 802.1-A:* Overview and Architecture
- *IEEE 802.1-B:* LAN/MAN Management
- *IEEE 802.1-D:* Transparent Bridges
- *IEEE 802.1-F:* Guidelines for the Development of Layer Management Standards
- *IEEE 802.10:* Secure Data Exchange

Note

At the time of this writing, key participants of the IEEE 802.11 standard effort included the following:

Victor Hayes, Chair

Stuart Kerry and Chris Zegelin, Vice Chairs

Bob O'Hara and Greg Ennis, Chief Technical Editors

George Fishel and Carolyn Heide, Secretaries

David Bagby, MAC Group Chair

Jan Boer, Direct Sequence Chair

Dean Kawaguchi, PHY Group and Frequency Hopping Chair

C. Thoman Baumgartner, Infrared Chair

HIPERLAN

High Performance Radio Local Area Network (HIPERLAN) is a European family of standards that specify high-speed digital wireless communication in the 5.15–5.3

GHz and the 17.1–17.3 GHz spectrum. These standards specify the Physical and Data Link Layers of network architecture, similar in scope to 802.11. However,

continues

continued

HIPERLAN operates using different protocols and is not compatible with other IEEE standards, such as IEEE 802.2 Logical Link Control.

Two stations in a HIPERLAN can exchange data directly, without any interaction from a wired network infrastructure. The simplest HIPERLAN consists of two stations. If two HIPERLAN stations are out of range with each other, a third station can relay the messages. HIPERLAN networks have the following specifications:

- Short range, approximately 150 feet (50 meters)
- Support of asynchronous and isochronous traffic
- Support of audio at 32 Kbps
- Support of video at 2 Mbps
- Support of data at 10 Mbps

HIPERLAN is unlikely to be a serious competitor to 802.11-based LANs, especially outside of Europe.

IEEE 802.11 Topology

The IEEE 802.11 topology consists of components, interacting to provide a wireless LAN that enables station mobility transparent to higher protocol layers, such as the LLC. A station is any device that contains functionality of the 802.11 protocol (that is, MAC Layer, PHY Layer, and interface to a wireless medium). The functions of the 802.11 standard reside physically in a radio NIC, the software interface that drives the NIC, and access point. The 802.11 standard supports the following two topologies:

- Independent Basic Service Set (IBSS) networks
- Extended Service Set (ESS) networks

These networks utilize a basic building block the 802.11 standard refers to as a BSS, providing a coverage area whereby stations of the BSS remain fully connected. A station is free to move within the BSS, but it can no longer communicate directly with other stations if it leaves the BSS.

Note

Harris Semiconductor was the first company to offer a complete radio chip set (called PRISM) for direct sequence spread spectrum that is fully compliant with IEEE 802.11. The PRISM chip set includes six integrated microcircuits that handle all signal processing requirements of 802.11.

Independent Basic Service Set (IBSS) Networks

An IBSS is a stand-alone BSS that has no backbone infrastructure and consists of at least two wireless stations (see Figure 3.6). This type of network is often referred to as an *ad hoc network* because it can be constructed quickly without much planning.

The ad hoc wireless network will satisfy most needs of users occupying a smaller area, such as a single room, a sales floor, or a hospital wing.




FIGURE 3.6 An independent BSS (IBSS) is the most basic type of 802.11 wireless LAN.

Extended Service Set (ESS) Networks

For requirements exceeding the range limitations of an independent BSS, 802.11 defines an Extended Service Set (ESS) LAN, as illustrated in Figure 3.7. This type of configuration satisfies the needs of large-coverage networks of arbitrary size and complexity.




FIGURE 3.7 An Extended Service Set (ESS) 802.11 wireless LAN consists of multiple cells interconnected by access points and a distribution system, such as ethernet.

The 802.11 standard recognizes the following mobility types:

- *No-transition*: This type of mobility refers to stations that do not move and those that are moving within a local BSS.
- *BSS-transition*: This type of mobility refers to stations that move from one BSS in one ESS to another BSS within the same ESS.
- *ESS-transition*: This type of mobility refers to stations that move from a BSS in one ESS to a BSS in a different ESS.

The 802.11 standard clearly supports the no-transition and BSS-transition mobility types. The standard, however, does not guarantee that a connection will continue when making an ESS-transition.

The 802.11 standard defines the *distribution system* as an element that interconnects BSSs within the ESS via access points. The distribution system supports the 802.11 mobility types by providing logical services necessary to handle address-to-destination mapping and seamless integration of multiple BSSs. An *access point* is an addressable station, providing an interface to the distribution system for stations located within various BSSs. The independent BSS and ESS networks are transparent to the LLC Layer.

Within the ESS, the 802.11 standard accommodates the following physical configuration of BSSs:

- *BSSs that partially overlap*: This type of configuration provides contiguous coverage within a defined area, which is best if the application cannot tolerate a disruption of network service.
- *BSSs that are physically disjointed*: For this case, the configuration does not provide contiguous coverage. 802.11 does not specify a limit to the distance between BSSs.
- *BSSs that are physically collocated*: This may be necessary to provide a redundant or higher-performing network.

The 802.11 standard does not constrain the composition of the distribution system; therefore, it may be 802-compliant or some nonstandard network. If data frames need transmission to and from a non-IEEE 802.11 LAN, these frames, as defined by the 802.11 standard, enter and exit through a logical point called a *portal*. The portal provides logical integration between existing wired LANs and 802.11 LANs. When the distribution system is constructed with 802-type components, such as 802.3 (ethernet) or 802.5 (token ring), the portal and the access point become one and the same.

IEEE 802.11 Logical Architecture

A topology provides a means of explaining necessary physical components of a network, but the *logical architecture* defines the network's operation. As Figure 3.8 illustrates, the logical architecture of the 802.11 standard that applies to each station consists of a single MAC and one of multiple PHYs.

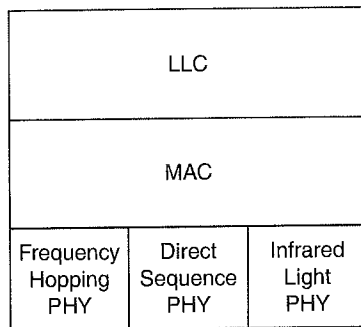


FIGURE 3.8 A single 802.11 MAC Layer supports three separate PHYs: frequency hopping spread spectrum, direct sequence spread spectrum, and infrared light.

IEEE 802.11 MAC Layer

The goal of the MAC Layer is to provide access control functions (such as addressing, access coordination, frame check sequence generation and checking, and LLC PDU delimiting) for shared-medium PHYs in support of the LLC Layer. The MAC Layer performs the addressing and recognition of frames in support of the LLC. The 802.11 standard uses CSMA/CA (carrier sense multiple access with collision avoidance); whereas, standard ethernet uses CSMA/CD (carrier sense multiple access with collision detection). It is not possible to both transmit and receive on the same channel using radio transceivers; therefore, an 802.11 wireless LAN takes measures only to avoid collisions, not to detect them.

IEEE 802.11 Physical Layers

The working group decided in July 1992 to concentrate its radio frequency studies and standardization efforts on the 2.4 GHz spread spectrum ISM bands for both the direct sequence and frequency hopping PHYs. The final standard specifies 2.4 GHz because this band is available license free in most parts of the world. The FCC Part 15 in the United States governs the radiated RF power in the ISM bands. Part 15 limits antenna gain to 6 dBi maximum and radiated power to one watt within the United States. European and Japanese regulatory groups limit radiated power to 10 milliwatts per 1 MHz. The actual frequencies authorized for use in the United States, Europe, and Japan differ slightly.

In March 1993, the 802.11 committee began receiving proposals for a direct sequence Physical Layer standard. After much discussion and debate, the committee agreed to include a chapter in the standard specifying the use of direct sequence. The direct sequence Physical Layer specifies two data rates:

- 2 Mbps using Differential Quaternary Phase Shift Keying (DQPSK) modulation
- 1 Mbps using Differential Binary Phase Shift Keying (DBPSK)

The standard defines seven direct sequence channels. One channel is exclusively available for Japan. Three channel pairs are defined for the United States and Europe. Channels in a pair can work without interference. In addition, the channels of all three pairs can be used simultaneously for redundancy or higher performance by developing a frequency plan that avoids signal conflicts.

In contrast to direct sequence, the 802.11-based frequency hopping PHY uses radios to send data signals by hopping from one frequency to another, transmitting a few bits on each frequency before shifting to a different one. Frequency hopping systems hop in a pattern that appears to be random, but really has a known sequence. A particular hop sequence is commonly referred to as a *frequency hopping channel*. Frequency hopping systems tend to be less costly to implement and do not consume as much power as their direct sequence counterpart, making them more suitable for portable applications. However, frequency hopping is much less tolerant of multiple-path and other interference sources. The system must retransmit data if it becomes corrupted on one of the hop sequence frequencies.

The 802.11 committee defined the frequency hopping Physical Layer to have a 1 Mbps data rate using 2-level Gaussian frequency shift keying (GFSK). This specification describes 79 channel center frequencies identified for the United States, from which there are three sets of 22 hopping sequences defined.

The infrared Physical Layer describes a modulation type that operates in the 850 to 950 nM band for small equipment and low-speed applications. The basic data rate of this infrared medium is 1 Mbps using 16-PPM (pulse position modulation) and an enhanced rate of 2 Mbps using 4-PPM. Peak power of infrared-based devices are limited to a peak power of 2 watts.

As with the IEEE 802.3 standard, the 802.11 Working Group is considering additional PHYs as applicable technologies become available.

For an inside look of each layer of the 802.11 standard, refer to Chapter 4, "Medium Access Control (MAC) Layer," and Chapter 5, "Physical (PHY) Layer."

APPENDIX B

IEEE 802.11 Services

The 802.11 standard defines *services* that provide the functions that the LLC Layer requires for sending MSDUs (MAC service data units) between two entities on the network. These services, which the MAC Layer implements, fall into two categories:

- Station services
 - Authentication
 - Deauthentication
 - Privacy
 - MSDU delivery
- Distribution system services
 - Association
 - Disassociation
 - Distribution
 - Integration
 - Reassociation

The following sections define the station and distribution system services.

Station Services

The 802.11 standard defines services for providing functions among stations. A station may be within any wireless element on the network, such as a handheld PC or handheld scanner. In addition, all access points implement station services. To provide necessary functionality, these stations need to send and receive MSDUs and implement adequate levels of security.

Authentication

Because wireless LANs have limited physical security to prevent unauthorized access, 802.11 defines authentication services to control LAN access to a level equal to a wired link. All 802.11 stations, whether they are part of an independent BSS or ESS network, must use the authentication service prior to establishing a connection (referred to as an association in 802.11 terms) with another station with which they will communicate. Stations performing authentication send a unicast management authentication frame to the corresponding station.

The IEEE 802.11 standard defines the following two authentication services:

- *Open system authentication*: This is the 802.11 default authentication method, which is a very simple, two-step process. First the station wanting to

authenticate with another station sends an authentication management frame containing the sending station's identity. The receiving station then sends back a frame alerting whether it recognizes the identity of the authenticating station.

- *Shared key authentication:* This type of authentication assumes that each station has received a secret shared key through a secure channel independent from the 802.11 network. Stations authenticate through shared knowledge of the secret key. Use of shared key authentication requires implementation of the Wireless Equivalent Privacy algorithm.

Deauthentication

When a station wishes to *disassociate* with another station, it invokes the *deauthentication* service. Deauthentication is a notification, and cannot be refused. Stations perform deauthentication by sending an authentication management frame (or group of frames to multiple stations) to *advise* the termination of authentication.

Privacy

With a wireless network, all stations and other devices can “hear” data traffic taking place within range on the network, seriously impacting the security level of a wireless link. IEEE 802.11 counters this problem by offering a privacy service option that raises the security level of the 802.11 network to that of a wired network.

The privacy service, applying to all data frames and some authentication management frames, is based on the 802.11 *Wired Equivalent Privacy (WEP)* algorithm that significantly reduces risks if someone eavesdrops on the network. This algorithm performs encryption of messages, as shown in Figure 3.9. With WEP, all stations initially start “in the clear”—that is, unencrypted. Refer to Chapter 4, in the section titled “Private Frame Transmissions,” for a description of how WEP works.




FIGURE 3.9 The Wired Equivalent Privacy (WEP) algorithm produces ciphertext, keeping eavesdroppers from “listening in” on data transmissions.

Note

The WEP protects RF data transmissions using a 64-bit seed key and the RC4 encryption algorithm. When enabled, WEC only protects the data packet information. Physical Layer headers are left unencrypted so that all stations can properly receive control information for managing the network.

Distribution System Services

Distribution system services, as defined by 802.11, provide functionality across a distribution system. Access points provide distribution system services. The following sections provide an overview of the services that distribution systems need to provide proper transfer of MSDUs.

Association

Each station must initially invoke the *association service* with an access point before it can send information through a distribution system. The association maps a station to the distribution system via an access point. Each station can associate with only a single access point, but each access point can associate with multiple stations. Association is also a first step to providing the capability for a station to be mobile between BSSs.

Disassociation

A station or access point may invoke the *disassociation service* to terminate an existing association. This service is a notification; therefore, neither party may refuse termination. Stations should disassociate when leaving the network. An access point, for example, may disassociate all its stations if being removed for maintenance.

Distribution

A station uses the *distribution service* every time it sends MAC frames across a distribution system. The 802.11 standard does not specify how the distribution system delivers the data. The distribution service provides the distribution system with only enough information to determine the proper destination BSS.

Integration

The *integration service* enables the delivery of MAC frames through a portal between a distribution system and a non-802.11 LAN. The integration function performs all required media or address space translations. The details of an integration function depends on the distribution system implementation and are beyond the scope of the 802.11 standard.

Reassociation

The *reassociation service* enables a station to change its current state of association. Reassociation provides additional functionality to support BSS-transition mobility for associated stations. The reassociation service enables a station to transition its association from one access point to another. This keeps the distribution system informed of the current mapping between access point and station as the station moves from BSS to BSS within an ESS. Reassociation also enables changing association attributes of an established association while the station remains associated with the same access point. The mobile station always initiates the reassociation service.

Note

IEEE 802.11 allows a client to roam among multiple access points that may be operating on the same or separate channels. To support the roaming function, each access point typically transmits a beacon signal every 100 milliseconds. Roaming stations use the beacon to gauge the strength of their existing access point connection. If the station senses a weak signal, the roaming station can implement the reassociation service to connect to an access point emitting a stronger signal.

**Case Study 3.5:
Reassociation Provides Roaming**

A grocery store in Gulfport, Mississippi, has a bar code–based shelf inventory system that helps the owners of the store keep track of what to stock, order, and so on. Several of the store clerks use handheld scanners during the store's closed hours to perform inventory functions. The store has a multiple cell 802.11-compliant wireless LAN (that is, ESS) consisting of access points A and B interconnected by an ethernet network. These two access points are sufficient to cover the store's entire floorspace and backroom.

At one end of the store in the frozen meat section, a clerk using a handheld device may associate with access point A. As the person walks with the device to the beer-and-wine section on the other end of the store, the mobile scanner (that is, the 802.11 station within the scanner) will begin sensing a signal from access point B. As the signal from B becomes stronger, the station will then *reassociate* with access point B, offering a much better signal for transmitting MSDUs.

Station States and Corresponding Frame Types

The state existing between a source and destination station (see Figure 3.10) governs which IEEE 802.11 frame types the two stations can exchange.

The following types of functions can occur within each class of frame:

- Class 1 Frames
 - Control Frames
 - Request to send (RTS)
 - Clear to send (CTS)
 - Acknowledgment (ACK)
 - Contention-free (CF)
 - Management Frames
 - Probe request/response
 - Beacon
 - Authentication




FIGURE 3.10 The operation of a station depends on its particular state.

Deauthentication

Announcement traffic indication message (ATIM)

- Data Frames
- Class 2 Frames
 - Management Frames
 - Association request/response
 - Reassociation request/response
 - Disassociation
- Class 3 Frames
 - Data Frames
 - Management Frames
 - Deauthentication
 - Control Frames
 - Power Save Poll

To keep track of station state, each station maintains the following two state variables:

- *Authentication state*: Has values of either unauthenticated and authenticated
- *Association state*: Has values of either unassociated and associated

Implications of the IEEE 802.11 Standard

As with any technologies and standards, one must be aware of the implications surrounding the implementation of wireless networks based on the IEEE 802.11 standard. Chapter 1, in the section “Wireless Network Concerns,” discusses the general issues of implementing wireless networks. In addition to these problems, the following are a couple of implications specifically related to the IEEE 802.11 standard:

- *Relatively low data rates:* As mentioned before, the 802.11 standard currently supports data rates up to 2 Mbps. Some end users and vendors claim this data rate is too low. In some cases, this is true; but in other cases, it is not true. Video transmissions, for example, may require higher data rates if applications need frame rates, pixel depth, and resolutions that require greater amounts of bandwidth. Large data block transmissions may also require higher data rates to keep transmission delays tolerable.

On the other hand, bar code applications, such as receiving, inventory, and price marking, generally work well under the 2 Mbps limitation of the current 802.11 standard.

- *Lack of standard roaming across multiple-vendor access points:* The 802.11 standard does not define the protocols necessary to move 802.11 frames within the distribution system because it falls outside the scope of 802-type LANs. The Network and Transport Layers are left to address distribution system protocols. As a result, the 802.11 standard does not define communications *between* access points.

Currently, it is up to the access point vendors to define the protocols necessary to support roaming from one access point to another. To be safe, you should consider purchasing access points from a single vendor, although you can mix-and-match radio cards in the appliances. Chapter 6, “Wireless System Integration,” discusses industry standards, such as the Inter Access Point Protocol (IAPP) specification, that are beginning to define multiple-vendor roaming protocols.

IEEE 802.11 Standard Compliance

No standard is worthwhile unless vendors and end users comply with it. The following sections describe activities taking place to ensure compliance with 802.11.

Vendor Compliance

Most wireless LAN vendors (that is, manufacturers of the hardware) are releasing initial radio cards and access points throughout 1998 and 1999 that comply with the official 802.11 standard. Before deeming their devices as 802.11 compliant, they must follow the protocol implementation compliance procedures that the 802.11 standard specifies in its appendix. The procedures state that the vendor shall

complete a Protocol Implementation Conformance Statement (PICS) proforma. The structure of the PICS proforma mainly includes a list of fixed questions that the vendor responds to with yes or no answers, indicating adherence to the standard. The PICS can have the following uses:

- A checklist that helps the vendor reduce the risk of failure to conform to the standard
- For the vendor and system implementor to better understand what 802.11-compliance means
- As a basis for designing an interface between the 802.11 device and another network or system
- As the basis for developing protocol conformance tests and simulations

To ensure proper compliance, vendors test their products at the InterOperability Laboratory located at the Leavitt Center on the campus of The University of New Hampshire. In March 1997, for example, Aironet Wireless Communications, Inc.; Breezecom Wireless Communications; Netwave Technologies, Inc.; Proxim Inc.; Raytheon Electronics; and Symbol Technologies performed joint interoperability testing to advance customer adoption of wireless technology. In some cases, users can upgrade their existing proprietary radio cards to be 802.11 compliant by just reinstalling NIC interface software on their appliances.

Note

Vendors are easing the transition to 802.11-compliant radio networks by offering relatively simple ways to upgrade existing radio LAN devices. Symbol, Inc., for example, offers a firmware upgrade to your existing Symbol 2.4 GHz (Spectrum 24) networks, avoiding the purchase of new network adapters.

The InterOperability Laboratory, founded in 1988, performs research and development work and is used by more than 100 vendors to verify the interoperability and conformance of their computer communications products. The University of New Hampshire encourages vendors to conduct interoperability testing by providing facilities for a multiple-vendor test environment. The goal of the laboratory is to provide complete testing for all networking products, including ethernet, ADSL, ATM, fast ethernet, FDDI, FDSE, Fibre Channel, gigabit ethernet, IP/Routing, Network Management, and Wireless.

Note

Be aware that in 1997 some vendors released "802.11-compliant" wireless LAN radio cards and access points that were not certified as compliant with the final 802.11 standard. These products may or may not operate within the final official standard.

WLI Forum

The Wireless LAN Interoperability Forum (WLI Forum), a not-for-profit corporation founded in March 1996, promotes the growth of the wireless LAN market by delivering interoperable products and services. The Forum consists primarily of appliance suppliers/vendors (such as Hewlett-Packard, Fujitsu, Monarch Marking Systems, and Handheld Products) having products that operate on the WLI Forum's OpenAir™ wireless network. The Forum provides certification via an independent third-party test lab to ensure proper compliance.

The OpenAir™ specification describes a MAC and radio frequency Physical Layer, similar in scope to the 802.11 specification. The OpenAir™ network is based on Proxim's RangeLAN2 protocol, employing frequency hopping spread spectrum technology in the unlicensed 2.4 GHz ISM band. The OpenAir™ operates at a data rate of 1.6 Mbps per channel, with 15 independent channels (hopping patterns) available. This architecture enables up to 15 wireless LANs to overlap independently in the same physical space,

providing up to 24 Mbps of aggregate network bandwidth.

The WLI Forum wrote the OpenAir™ specification to motivate third-party development of compatible products. At the time, with no official IEEE standard on wireless networking, the Forum decided to base its specification on Proxim's product. Soon after the release of the 802.11 specification in June 1997, the WLI Forum announced its support for the adoption of the IEEE 802.11 standard and urged the supplier community to move toward conformance. As a result, the WLI Forum is likely to establish conformity to the IEEE 802.11 standard as well.

The WLI Forum is a worldwide organization, and is completely self funded through membership dues and fees. Membership is open to all companies that develop, manufacture, or sell wireless LAN products or services. For more information on the WLI Forum, visit their Web site located at <http://www.wlif.com>.

End-User Compliance

Throughout 1999 and beyond, end users should begin widespread implementations of 802.11-compliant LANs. As an end user, do you need to purchase and use products that comply with the 802.11 standard? Of course the answer is no, but you should carefully consider the advantages and disadvantages of implementing 802.11-compliant networks. Most likely, complying with 802.11 will be favored over the use of proprietary networks unless extenuating circumstances prevail. If the decision is to go with 802.11, you will be starting with one of the following scenarios:

- No existing implementation of wireless LANs
- Existing implementation of proprietary wireless LANs

If you are an end user with no existing installation of wireless networking components, compliance with the 802.11 standard is easy. Right? Actually, it is not as sim-

APPENDIX B

ple as it seems. The 802.11 standard is not as Plug and Play as the 802.3 ethernet standard. With 802.11, you must first decide which version of 802.11 best satisfies your needs. You might consider the following questions:

- *What type of modulation do I need?* Do I have radio interference implications that lean toward using the infrared PHY? Does the application require wider area coverage that may depend on the longer range capability of one of the spread spectrum PHYs? If the choice is spread spectrum, should I use direct sequence or frequency hopping?
- *Will the application require roaming across BSS cells interconnected by access points of different vendors?* If yes, you will need to think about how to provide roaming between access points.
- *Does the network require the optional WEP security?* If the answer is yes, be sure to choose wireless devices having WEP available.
- *Do the appliances I need to comply with have the 802.11 options I have chosen?* If not, you need to choose options that comply with the appliance, or you must choose different appliances.

Answers to the preceding questions define the options you need to consider when planning to purchase radio cards and access points complying with 802.11.

If proprietary wireless LANs already exist, you will need to either upgrade or replace the existing network to make it compliant with 802.11. Many of the vendors offer free upgrades to make your existing wireless LANs (if they are of a recent enough version) compliant with 802.11. BreezeCOM, for example, guarantees software upgrades to the IEEE 802.11 standard for its BreezeNET PRO product line.

If it is not possible or feasible to upgrade your existing wireless LAN, then of course you must perform a complete replacement if benefits outweigh the expenses. The replacement of the network will be difficult to cost-justify; however, it may become necessary as proprietary wireless components become obsolete.

International Electromagnetic Compliance

The 802.11 standard specifies operation in the 2.4 GHz band; however, electromagnetic compatibility requirements vary from one country to another. Operating frequencies, power levels, and spurious levels differ throughout the world.

Regional and national regulatory administrations of each individual country demand certification of wireless equipment. The 802.11 standard, however, identifies the minimum technical requirements for interoperability and compliance based on established regulations for Europe, Japan, and the North America. Therefore, wireless LAN vendors must be aware of all current regulatory requirements prior to releasing a product for sale in a particular country. The following agencies and

APPENDIX B

PART I WIRELESS NETWORKS—A FIRST LOOK

122

documents specify the current regulatory requirements for various geographical areas:

Canada

- *Approval standards:* Industry Canada (IC)
- *Documents:* GL36
- *Approval authority:* Industry Canada

Europe

- *Approval standards:* European Telecommunications Standards Institute
- *Documents:* ETS 300-328, ETS 300-339
- *Approval authority:* National Type Approval Authorities

France

- *Approval standards:* La Reglementation en France por les Equipements fonctionnant dans la bande de frequences 2,4 GHz "RLAN-Radio Local Area Network"
- *Documents:* SP/DGPT/ATAS/23, ETS 300-328, ETS 300-339
- *Approval authority:* Direction Generale des Postes et Telecommunications

Japan

- *Approval standards:* Research and Development Center for Radio Communications (RCR)
- *Documents:* RCR STD-33A
- *Approval authority:* Ministry of Telecommunications (MKK)

Spain

- *Approval standards:* Suplemento del Numero 164 del Boletin Oficial del Estado (published 10 July 91; revised 25 June 93)
- *Documents:* ETS 300-328, ETS 300-339
- *Approval authority:* Cuadro Nacional De Atribucion De Frecuencias

The United States of America

- *Approval standards:* Federal Communications Commission (FCC)
- *Documents:* CFR47, Part 15, Sections 15.205, 15.209, 15.247
- *Approval authority:* FCC

Operation in countries within Europe and other areas outside Japan or North America may be subject to additional regulations.

IEEE 802.11 Working Group Operations

The 802.11 Working Group is a part of the IEEE LAN MAN Standards Committee (LMSC), which reports to the Standards Activity Board (SAB) of the IEEE Computer Society. IEEE 802.11 meetings are open to anyone. The only requirement to attend is to pay dues, which offset meeting expenses. Most of the active participants are representatives from companies developing wireless LAN components. The IEEE bylaws explain that to vote on standards activities, however, you must become a member by participating in at least two out of four consecutive plenary meetings. Then, you must continue to attend meetings to maintain voting status. The 802.11 Working Group meets three times a year during the plenary sessions of the IEEE 802 and three times a year between plenary sessions.

The IEEE 802.11 Working Group consists of about 200 members; membership falls into the following categories:

- *Voting members:* Those who have maintained voting status.
- *Nearly members:* Those who have participated in two sessions of meetings, one of which being a plenary session. Nearly members become voting members in the first session they attend following their qualification for nearly membership.
- *Aspirant members:* Those who have participated in one plenary or interim session meeting.
- *Sleeping voting members:* Those who were once voting members, but have chosen to discontinue.

Future of the IEEE 802.11 Standard

What is the future of IEEE 802.11? Will end users eventually fully comply with the standard? Will the 802.11 Working Group solve implications revolving around the standard? Only time will tell for certain. It is known today, however, that all major wireless LAN vendors are releasing 802.11-compliant wireless LANs throughout 1998, and these vendors are making it fairly easy for end users to upgrade their existing systems. This, combined with the advantages of standardization, should proliferate the use of 802.11-compliant networks.

To solve implications of the current release of the standard, the IEEE 802.11 Working Group is actively working on the following projects that will aid the widespread acceptance of the standard:

- *802.11rev: Revision of IEEE Standard 802.11-1997:* This project was charted to rectify a number of errors in the current standard and to accommodate input from the JTC1 review to result in a single JTC1/IEEE standard.

- *802.11a: Extension of the IEEE Standard 802.11-1997 with a higher data rate PHY in the 5 GHz band:* This project was initiated to develop a high speed (about 20 Mbps) wireless PHY suitable for data, voice, and image information services in fixed, moving, or portable wireless local area networks. The project concentrates on improving spectrum efficiency and will review the existing 802.11 MAC to ensure its capability to operate at the higher speeds.

The IEEE 802.11 Working Group will actively correspond with regulatory bodies worldwide to encourage spectrum allocations that match these frequencies.

- *802.11b: Extension of the IEEE Standard 802.11-1997 with a higher data rate PHY in the 2.4 GHz band:* The purpose of this project is to extend the performance and the range of applications of the existing 802.11 standard. The header of the two existing radio-based PHYs can support data rates up to 4.5 Mbps for frequency hopping and up to 25.5 Mbps for direct sequence. This project will investigate ways to exploit these data rate capabilities and analyze the capability of the existing 802.11 MAC to support higher data rates.

The actual data rates targeted by this project are at least 3 Mbps for the frequency hopping PHY and at least 8 Mbps for the direct sequence PHY. As with project 802.11a, IEEE 802.11 will correspond with regulatory bodies worldwide to ensure that the proposed extension will be applicable as widely as possible.

In addition to the preceding official projects, the 802.11 Working Group is actively studying the needs for standardization of wireless communications of wearable computing devices. The study is examining the requirements for Wireless Personal Area Networking (WPAN) of devices that are worn or carried by individuals. The objectives of the study group are as follows:

- Review WPAN requirements.
- Determine the need for a standard.
- If a standard is necessary, draft a PAR for submittal.
- Seek appropriate sponsorship within 802.

The study group is soliciting industry input on market requirements and technical solutions for a WPAN with 0-to-30-foot range, data rates of less than 1 Mbps, low power consumption, small size (less than 0.5 cubic inches), and low cost relative to target device.

As mentioned in this chapter, the 802.11 wireless LAN standard certainly has benefits that an organization should consider when selecting components that provide LAN mobility. IEEE 802 is a solid family of standards that will provide much greater multiple-level interoperability than proprietary systems.

APPENDIX B

Wireless LANs conforming to 802.11 provide interoperability between radio cards and access points. The 802.11 standard has the backing of IEEE, having an excellent track record of developing long-lasting standards, such as IEEE 802.3 (ethernet) and IEEE 802.5 (token ring). When designing a wireless LAN, definitely consider the use of 802.11-compliant products, but ensure that the data rates of 802.11 will support your application and that the chosen components support roaming between access points.

With 802.11, system implementors have several choices. You will need to choose the type of physical medium, for example: frequency hopping spread spectrum, direct sequence spread spectrum, or infrared light. This concept is similar to choosing between twisted-pair, optical-fiber, and coaxial cable in an ethernet LAN. You will also need to determine how to interface wireless devices with server operating systems and applications. In defining these elements, be sure the resulting network supports all requirements.

New Economy; Airborne and grass roots. By popular acclaim, a wireless format with a name only a geek could love is taking hold.

By John Markoff

Oct. 30, 2000

See the article in its original context from
October 30, 2000, Section C, Page 5 Buy Reprints

[VIEW ON TIMESMACHINE](#)

TimesMachine is an exclusive benefit for home
delivery and digital subscribers.

AT the recent Agenda 2001 computer conference in Phoenix earlier this month, there was validation and a touch of irony for the conference organizer, Robert Metcalfe, co-inventor of the Ethernet office networking standard.

Hundreds of those attending the conference sat in a huge auditorium with their portable computers wirelessly linked to the Internet via tiny PC cards plugged in to their machines.

It was tacit validation of a theorem Mr. Metcalfe set forth years ago, now widely known as Metcalfe's Law, which states that "the usefulness, or utility, of a network equals the square of the number of users."

The irony, of course, was that while many of the conference participants were using a technology that essentially provides wireless connections at Ethernet speeds, they were using it to read their e-mail and surf the Web rather than pay much attention to Mr. Metcalfe's conference.

There is no doubt, however, that "wireless Ethernet"-- formally known as the 802.11b wireless technical standard as specified by the Institute of Electrical and Electronics Engineers -- is finally taking off.

APPENDIX C

The Ethernet standard for wiring computers into local networks caught on two decades ago because it was "open" -- owned by no single company and available for many to adopt and improve. Those same characteristics could propel wireless Ethernet as embodied in the 802.11b standard, which allows data to be transmitted at 11 megabits -- 11 million bits a second.

The 802.11b format is catching on so quickly that it is displacing alternative wireless competitors that include Bluetooth and HomeRF. The cost of 802.11b technology continues to plummet; chip sets that cost as little as \$10 or less may arrive in the next two years. So it should become cheaper and easier to set up an office network wirelessly than with traditional Ethernet wires.

To be sure, there are some clouds ahead. The 802.11b wireless transmitters operate on the 2.4 gigahertz radio band, which does not require a license to use. Some technical experts worry that this band may soon grow so congested that it will create the world's first wireless data gridlock.

The standard, first popularized by Apple Computer in its Airport line of wireless products last year, is now being embraced so quickly that it is touching off a wireless "air rush" as start-up companies and telecommunication vendors vie to lock up valuable sites at airports, hotels and other public hot spots. Such companies operate the Internet server computers by which wireless users actually connect to the global network.

The appearance of the wireless standard in public spaces is following on the heels of installations on university campuses and corporate office parks. And some community 802.11b wireless networks have been set up, including SFlan in San Francisco.

As part of SFlan, some Internet hobbyists have set up inexpensive 802.11b networks on their rooftops and are distributing Internet service throughout their neighborhoods. One user, Tim Pozar, said his local network reached a half-mile radius around his home.

Brewster Kahle, a computer network expert who has led the SFlan project, said, "It's possible that a grass-roots broadband network could be built organically."

But for now, commercial efforts seem to be gathering steam most quickly. There have been a series of announcements in recent weeks by wireless companies including Aerzone , Mobilstar and Wayport, that have struck deals with airports and hotels to install 802.11b.

Just last Friday, United Airlines said it was teaming with Aerzone, a San Francisco-based subsidiary of Softnet Systems, to deploy 802.11b in Red Carpet Club airport lounges, gate areas and terminals in as many as 50 airports served by the airline.

APPENDIX C

"This is potentially a huge business because we offer the two things people want most: relatively unlimited bandwidth and mobility," said Lawrence B. Brilliant, chief executive of Aerzone.

Just two days earlier Wayport, based in Austin, Tex., announced it was installing 802.11b in the lobbies of 15 hotels in the Los Angeles area. Wayport has already started its service at the Dallas-Fort Worth International Airport and the Austin Bergstrom airport.

And while so far Apple and I.B.M. are the only two computer makers to offer portable systems with built-in 802.11b capability, by early next year the standard is expected to become a common built-in feature on all makes of portable computers.

Meanwhile, a number of sports stadium deals have been announced and several of the wireless start-up companies say they have been in talks with Starbucks to offer wireless Internet service in the company's nationwide chain of coffee shops.

Indeed the possibility of Starbucks's encouraging customers to spend time sipping coffee in its stores while they read their e-mail has created a new technology buzz-phrase: the "high-loiter retail" marketplace, in the words of Brett Stewart, Wayport's president and founder.

The wireless networks based on 802.11b are also becoming popular as a convenient and low-cost way to create a network within homes already connected to the Internet through D.S.L. or cable modems.

All this activity raises the possibility that 802.11b might upset the plans of some of the big telecommunications giants that are planning to spend millions of dollars building third-generation data and voice cellular networks. If millions of computer users and companies effectively build their own high-speed data network from the ground up, the telecommunications carriers might think twice about putting money into third-generation systems.

The 802.11b standards offer far greater speed than the proposed third-generation network standards, which generally offer two megabits that must be shared by all the users of a single cell. And the industry is finishing a standard called 802.11a, which will allow even higher speed -- 54 megabits a second -- on the 5 gigahertz radio band.

In fact, a top Microsoft executive, Craig Mundie, said his company was trying to rally the computer and telecommunications industries to agree upon that standard for the future of wireless data networking.

If the largely spontaneous 802.11 wave does swamp the various other wireless data networking standards, it might be fitting. After all, it was as an anarchic self-assembling world of isolated networks that the Internet originally came into being.

3/16/2021

New Economy; Airborne and grass roots. By popular acclaim, a wireless format with a name only a geek could love is taking hold. - The ...

A version of this article appears in print on , Section C, Page 5 of the National edition with the headline: New Economy; Airborne and grass roots. By popular acclaim, a wireless format with a name only a geek could love is taking hold.

APPENDIX C

Abstract

A few years ago it was recognized that the vision of a truly low-cost, low-power radio-based cable replacement was feasible. Such a ubiquitous link would provide the basis for portable devices to communicate together in an ad hoc fashion by creating personal area networks which have similar advantages to their office environment counterpart, the local area network. Bluetooth™ is an effort by a consortium of companies to design a royalty-free technology specification enabling this vision. This article describes the radio system behind the Bluetooth concept. Designing an ad hoc radio system for worldwide usage poses several challenges. The article describes the critical system characteristics and motivates the design choices that have been made.

The Bluetooth Radio System

Jaap C. Haartsen, Ericsson Radio Systems B.V.

In the last decades, progress in microelectronics and very large scale integration (VLSI) technology has fostered the widespread use of computing and communication devices for commercial usage. The success of consumer products like PCs, laptops, personal digital assistants (PDAs), cell phones, cordless phones, and their peripherals has been based on continuous cost and size reduction. Information transfer between these devices has been cumbersome, mainly relying on cables. Recently, a new universal radio interface has been developed enabling electronic devices to communicate wirelessly via short-range ad hoc radio connections. The Bluetooth technology — which has gained the support of leading manufacturers like Ericsson, Nokia, IBM, Toshiba, Intel, and many others — eliminates the need for wires, cables, and the corresponding connectors between cordless or mobile phones, modems, headsets, PDAs, computers, printers, projectors, and so on, and paves the way for new and completely different devices and applications. The technology enables the design of low-power, small-sized, low-cost radios that can be embedded in existing (portable) devices. Eventually, these embedded radios will lead toward ubiquitous connectivity and truly connect everything to everything. Radio technology will allow this connectivity to occur without any explicit user interaction.

This article describes the basic design and technology trade-offs which have led to the Bluetooth radio system. We describe some fundamental issues regarding ad hoc radio systems. We give an overview of the Bluetooth system itself with the emphasis on the radio architecture. It explains how the system has been optimized to support ad hoc connectivity. We also describe the Bluetooth specification effort.

Ad Hoc Radio Connectivity

The majority of radio systems in commercial use today are based on a cellular radio architecture. A mobile network established on a wired backbone infrastructure uses one or more base stations placed at strategic positions to provide local cell coverage; users apply portable phones, or more generic mobile terminals, to access the mobile network; the terminals maintain a connection to the network via a radio link to the base stations. There is a strict separation between the base stations and the terminals. Once registered to the network, the terminals remain locked to the control channels in the network, and connections can be established and released according to the control channel protocols. Channel access, channel allocation, traffic control, and interference minimization are neatly con-

trolled by the base stations. Examples of these conventional radio systems are the public cellular phone systems like Global System for Mobile Communications (GSM), D-AMPS, and IS-95 [1–3], but also private systems like wireless local area network (WLAN) systems based on 802.11 or HIPERLAN I and HIPERLAN II [4–6], and cordless systems like Digital Enhanced Cordless Telecommunications (DECT) and Personal Handyphone System (PHS) [7, 8].

In contrast, in truly ad hoc systems, there is no difference between radio units; that is, there are no distinctive base stations or terminals. Ad hoc connectivity is based on peer communications. There is no wired infrastructure to support connectivity between portable units; there is no central controller for the units to rely on for making interconnections; nor is there support for coordination of communications. In addition, there is no intervention of operators. For the scenarios envisioned by Bluetooth, it is highly likely that a large number of ad hoc connections will coexist in the same area without any mutual coordination; that is, tens of ad hoc links must share the same medium at the same location in an uncoordinated fashion. This is different from ad hoc scenarios considered in the past, where ad hoc connectivity focused on providing a single (or very few) network(s) between the units in range [4, 5]. For the Bluetooth applications, typically many independent networks overlap in the same area. This will be indicated as a scatter ad hoc environment. Scatter ad hoc environments consist of multiple networks, each containing only a limited number of units. The difference between a conventional cellular environment, a conventional ad hoc environment, and a scatter ad hoc environment is illustrated in Fig. 1. The environmental characteristics the ad hoc radio system has to operate in have a major impact on the following fundamental issues:

- Applied radio spectrum
- Determining which units are available to connect to
- Connection establishment
- Multiple access scheme
- Channel allocation
- Medium access control
- Service prioritization (i.e., voice before data)
- (Mutual) interference
- Power consumption

Ad hoc radio system have been in use for some time, for example, walky-talky systems used by the military, police, fire departments, and rescue teams in general. However, the Bluetooth system is the first commercial ad hoc radio system envisioned to be used on a large scale and widely available to the public.

Bluetooth Radio System Architecture

In this section the technical background of the Bluetooth radio system is presented. It describes the design trade-offs made in order to optimize the ad hoc functionality and addresses the issues listed above.

Radio Spectrum


The choice of radio spectrum is first determined by the lack of operator interaction. The spectrum must be open to the public without the need for licenses. Second, the spectrum must be available worldwide. The first Bluetooth applications are targeted at the traveling businessperson who connects his/her portable devices wherever he/she goes. Fortunately, there is an unlicensed radio band that is globally available. This band, the Industrial, Scientific, Medical (ISM) band, is centered around 2.45 GHz and was formerly reserved for some professional user groups but has recently been opened worldwide for commercial use. In the United States, the band ranges from 2400 to 2483.5 MHz, and the FCC Part 15 regulations apply. In most parts of Europe,¹ the same band is available under the ETS-300328 regulations. In Japan, recently the band from 2400 to 2500 MHz has been allowed for commercial applications and has been harmonized with the rest of the world. Summarizing, in most countries of the world, free spectrum is available from 2400 MHz to 2483.5 MHz, and harmonization efforts are ongoing to have this radio band available truly worldwide.

The regulations in different parts of the world differ. However, their scope is to enable fair access to the radio band by an arbitrary user. The regulations generally specify the spreading of transmitted signal energy and maximum allowable transmit power. For a system to operate globally, a radio concept has to be found that satisfies all regulations simultaneously. The result will therefore be the minimum denominator of all the requirements.

Interference Immunity

Since the radio band is free to be accessed by any radio transmitter as long as it satisfies the regulations, interference immunity is an important issue. The extent and nature of the interference in the 2.45 GHz ISM band cannot be predicted. Radio transmitters may range, for example, from 10 dBm baby monitors to 30 dBm WLAN access points. With high probability, the different systems sharing the same band will not be able to communicate. Coordination is therefore not possible. More of a problem are the high-power transmitters covered by the FCC part 18 rules which include, for example, microwave ovens and lighting devices. These devices fall outside the power and spreading regulations of part 15, but still coexist in the 2.45 GHz ISM band. In addition to interference from external sources, co-user interference must be taken into account, which results from other Bluetooth users.

Interference immunity can be obtained by interference suppression or avoidance. Suppression can be obtained by coding or direct-sequence spreading. However, the dynamic range of the interfering and intended signals in an ad hoc, uncoordinated radio environment can be huge. Taking into account the distance ratios and power differences of uncoordinated transmitters, near-far ratios in excess of 50 dB are no exception. With desired user rates on the order of 1 Mb/s and beyond, practically attained coding and processing gains are inadequate. Instead, interference avoidance is more attractive




■ **Figure 1.** Topologies for: a) cellular radio systems with squares representing stationary base stations; b) conventional ad hoc systems; and c) scatter ad hoc systems.

since the desired signal is transmitted at points in frequency and/or time where interference is low or absent. Avoidance in time can be an alternative if the interference concerns a pulsed jammer and the desired signal can be interrupted. Avoidance in frequency is more practical. Since the 2.45 GHz band provides about 80 MHz of bandwidth and most radio systems are band-limited, with high probability a part of the radio spectrum can be found where there is no dominant interference. Filtering in the frequency domain provides the suppression of the interferers at other parts of the radio band. The filter suppression can easily arrive at 50 dB or more.

Multiple Access Scheme

The selection of the multiple access scheme for ad hoc radio systems is driven by the lack of coordination and the regulations in the ISM band. Frequency-division multiple access (FDMA) is attractive for ad hoc systems since channel orthogonality only relies on the accuracy of the crystal oscillators in the radio units. Combined with an adaptive or dynamic channel allocation scheme, interference can be avoided. Unfortunately, pure FDMA does not fulfill the spreading requirements set in the ISM band. Time-division multiple access (TDMA) requires strict time synchronization for channel orthogonality. For multiple collocated ad hoc connections, maintaining a common timing reference becomes rather cumbersome. Code-division multiple access (CDMA) offers the best properties for ad hoc radio systems since it provides spreading and can deal with uncoordinated systems. Direct sequence (DS)-CDMA is less attractive because of the near-far problem which requires coordinated power control or excessive processing gain. In addition, as in TDMA, DS-CDMA channel orthogonality requires a common timing reference. Finally, for higher user rates, rather high chip rates are required, which is less attractive because of the wide bandwidth (interference immunity) and higher current consumption. Frequency-hopping (FH)-CDMA combines a number of properties which make it the best choice for ad hoc radio systems. On average the signal can be spread over a large frequency range, but instantaneously only a small bandwidth is occupied, avoiding most of the potential interference in the ISM band. The hop carriers are orthogonal, and the interference on adjacent hops can effectively be suppressed by filter-

¹ In France and Spain the exact location of the band differs, and the band is smaller.



■ **Figure 2.** An illustration of the FH/TDD channel applied in Bluetooth.

ing. The hop sequences will not be orthogonal (coordination of hop sequences is not allowed by the FCC rules anyway), but narrowband and co-user interference is experienced as short interruptions in the communications, which can be overcome with measures at higher-layer protocols.

Bluetooth is based on FH-CDMA. In the 2.45 GHz ISM band, a set of 79 hop carriers have been defined at a 1 MHz spacing.² The channel is a hopping channel with a nominal hop dwell time of 625 μ s. A large number of pseudo-random hopping sequences have been defined. The particular sequence is determined by the unit that controls the FH channel, which is called the *master*. The native clock of the master unit also defines the phase in the hopping sequence. All other participants on the hopping channel are *slaves*; they use the master identity to select the same hopping sequence and add time offsets to their respective native clocks to synchronize to the frequency hopping. In the time domain, the channel is divided into slots. The minimum dwell time of 625 μ s corresponds to a single slot. To simplify implementation, full-duplex communications is achieved by applying time-division duplex (TDD). This means that a unit alternately transmits and receives. Separation of transmission and reception in time effectively prevents crosstalk between the transmit and receive operations in the radio transceiver, which is essential if a one-chip implementation is desired. Since transmission and reception take place at different time slots, transmission and reception also take place at different hop carriers. Figure 2 illustrates the FH/TDD channel applied in Bluetooth. Note that multiple ad hoc links will make use of different hopping channels with different hopping sequences and have misaligned slot timing.

The Modulation Scheme

In the ISM band, the signal bandwidth of FH systems is limited to 1 MHz. For robustness, a binary modulation scheme was chosen. With the above-mentioned bandwidth restriction, the data rates are limited to about 1 Mb/s. For FH systems and support for bursty data traffic, a noncoherent detection scheme is most appropriate. Bluetooth uses Gaussian-shaped frequency shift keying (FSK) modulation with a nominal modulation index of $k = 0.3$. Logical ones are sent as positive frequency deviations, logical zeroes as negative frequency deviations. Demodulation can simply be accomplished by a limiting FM discriminator. This modulation scheme allows the implementation of low-cost radio units.

Medium Access Control

Bluetooth has been optimized to allow a large number of uncoordinated communications to take place in the same area. Unlike other ad hoc solutions where all units in range share the same channel, Bluetooth has been designed to allow

a large number of independent channels, each channel serving only a limited number of participants. With the considered modulation scheme, a single FH channel in the ISM band only supports a gross bit rate of 1 Mb/s. This capacity has to be shared by all participants on the channel. Theoretically, the spectrum with 79 carriers can support 79 Mb/s. In the user scenarios targeted by Bluetooth, it is highly unlikely that all units in range need to share information among all of them. By using a large number of independent 1 Mb/s channels to which only the units are connected that really want to exchange information, the 80 MHz is exploited much more effectively. Due to nonorthogonality of the hop sequences, the theoretical capacity of 79 Mb/s cannot be reached, but is at least much larger than 1 Mb/s.

An FH Bluetooth channel is associated with a piconet. As mentioned earlier, the piconet channel is defined by the identity (providing the hop sequence) and system clock (providing the hop phase) of a master unit. All other units participating in the piconet are slaves. Each Bluetooth radio unit has a free-running system or native clock. There is not a common timing reference, but when a piconet is established, the slaves add offsets to their native clocks to synchronize to the master. These offsets are released again when the piconet is cancelled, but can be stored for later use. Different channels have different masters and therefore also different hopping sequences and phases. The number of units that can participate on a common channel is deliberately limited to eight (one master and seven slaves) in order to keep a high-capacity link between all the units. It also limits the overhead required for addressing. Bluetooth is based on peer communications. The master/slave role is only attributed to a unit for the duration of the piconet. When the piconet is cancelled, the master and slave roles are cancelled. Each unit can become a master or slave. By definition, the unit that establishes the piconet becomes the master.

In addition to defining the piconet, the master also controls the traffic on the piconet and takes care of access control. Access is completely contention free. The short dwell time of 625 μ s only allows the transmission of a single packet. A contention-based access scheme would provide too much overhead and is not efficient in the short dwell time Bluetooth applies. In Bluetooth, the master implements centralized control; only communication between the master and one or more slaves is possible. The time slots are alternately used for master transmission and slave transmission. In master transmission, the master includes a slave address of the unit for which the information is intended. In order to prevent collisions on the channel due to multiple slave transmissions, the master applies a polling technique: for each slave-to-master slot, the master decides which slave is allowed to transmit. This decision is performed on a per-slot basis: only the slave addressed in the master-to-slave slot directly preceding the slave-to-master slot is allowed to transmit in the slave-to-master slot. If the master has information to send to a specific slave, this slave is polled implicitly and can return information. If the master has no information to send, it has to poll the slave explicitly with a short poll packet. Since the master schedules the traffic in both the uplink and downlink, intelligent scheduling algorithms have to be used that take into account the slave characteristics. The master control effectively prevents collisions between the participants on the piconet channel. Independent collocated piconets may interfere when they occasionally use the same hop carrier. A type of ALOHA is applied: information is transmitted without checking for a clear carrier (listen-before-talk). If the information is received incorrectly, it is retransmitted at

² Currently, for France and Spain a reduced set of 23 hop carriers has been defined at a 1 MHz carrier spacing.


the next transmission opportunity (for data only). Due to the short dwell time, collision avoidance schemes are less appropriate for FH radio. For each hop, different contenders are encountered. Backoff mechanisms are therefore less efficient.

Packet-Based Communications

The Bluetooth system uses packet-based transmission: the information stream is fragmented into packets. In each slot, only a single packet can be sent. All packets have the same format, starting with an access code, followed by a packet header, and ending with the user payload (Fig. 3).

The access code has pseudo-random properties and is used as a direct-sequence code in certain access operations. The access code includes the identity of the piconet master. All packets exchanged on the channel are identified by this master identity. Only if the access code matches the access code corresponding to the piconet master will the packet be accepted by the recipient. This prevents packets sent in one piconet falsely being accepted by units of another piconet that happens to land on the same hop carrier. In the receiver, the access code is matched against the anticipated code in a sliding correlator. This correlator provides the direct-sequence processing gain. The packet header contains link control information: a 3-bit slave address to separate the slaves on the piconet, a 1-bit acknowledgment/negative acknowledgment (ACK/NACK) for the automatic repeat request (ARQ) scheme, a 4-bit packet type code to define 16 different payload types, and an 8-bit header error check (HEC) code which is a cyclic redundancy check (CRC) code to detect errors in the header. The packet header is limited to 18 information bits in order to restrict the overhead. The header is further protected by 1/3 rate forward error correction (FEC) coding. Bluetooth defines four control packets:

- The ID or identification packet: Only consists of the access code; used for signaling
- The NULL packet: Only has an access code and a packet header; used if link control information carried by the packet header has to be conveyed
- The POLL packet: Similar to the NULL packet; used by the master to force slaves to return a response
- The FHS packet: An FH-synchronization packet; used to exchange real-time clock and identity information between the units; contains all the information to get two units hop synchronized



■ Figure 3. The format of packets applied in Bluetooth.

The remaining 12 type codes are used to define packets for synchronous and asynchronous services. These 12 types are divided into three segments. Segment 1 specifies packets that fit into a single slot, segment 2 specifies 3-slot packets, and segment 3 specifies 5-slot packets. Multislot packets are sent on a single-hop carrier. The hop carrier which is valid in the first slot is used for the remainder of the packet; therefore, there is no frequency switch in the middle of a packet. After the packet has been sent, the hop carrier as specified by the current master clock value is used (Fig. 4). Note that only an odd number of multislot packets have been defined, which guarantees that the TX/RX timing is maintained.


On the slotted channel, synchronous and asynchronous links have been defined, as will be further explained later. The interpretation of packet type is different for synchronous and asynchronous links. Currently, asynchronous links support payloads with or without a 2/3-rate FEC coding scheme. In addition, on these links single-slot, three-slot, and five-slot packets are available. The maximum user rate that can be obtained over the asynchronous link is 723.2 kb/s. In that case, a return link of 57.6 kb/s can still be supported. Link adaptation can be applied on the asynchronous link by changing the packet length and FEC coding depending on link conditions. The payload length is variable and depends on the available user data. However, the maximum length is limited by the minimum switching time between RX and TX, which is specified at 200 μ s. This switching time seems large, but allows the use of open-loop voltage controlled oscillators (VCOs) for direct modulation and provides time for packet processing between RX and TX; this is also discussed later. For synchronous links, only single-slot packets have been defined. The payload length is fixed. Payloads with 1/3-rate FEC, 2/3-rate, or no FEC are supported. The synchronous link supports a full-duplex link with a user rate of 64 kb/s in both directions.

Physical Link Definition

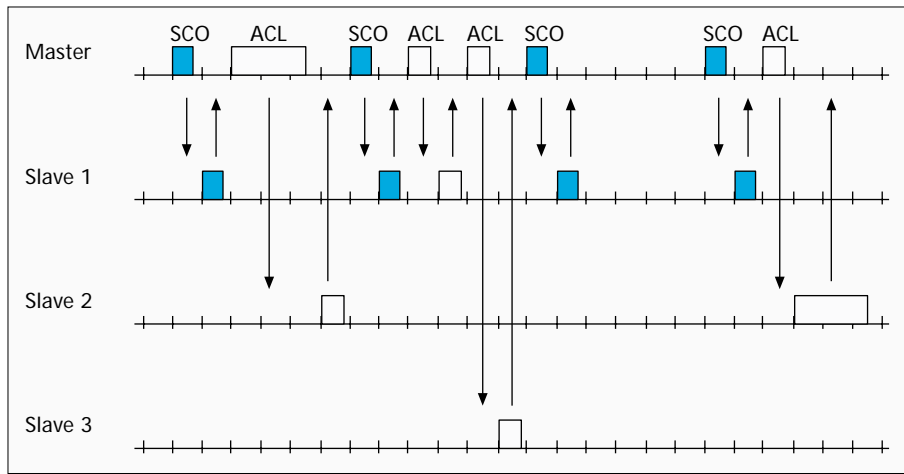
The Bluetooth link supports both synchronous services such as voice traffic, and asynchronous services such as bursty data traffic. Two physical link types have been defined:

- The synchronous connection-oriented (SCO) link
- The asynchronous connectionless (ACL) link


The SCO link is a point-to-point link between the master and a single slave. The link is established by reservation of duplex slots at regular intervals. The ACL link is a point-to-multipoint link between the master and all the slaves on the piconet. The ACL link can use all of the remaining slots on the channel not used for SCO links. The traffic over the ACL link is scheduled by the master. The slotted structure of the piconet channel allows effective mixing of the synchronous and asynchronous links. An example of a channel with SCO and ACL links is



■ Figure 4. The frequency and timing characteristics of single-slot, three-slot, and five-slot packets.



■ **Figure 5.** An example of mixing synchronous SCO links and asynchronous ACL links on a single piconet channel.



■ **Figure 6.** Frequency and timing behavior for a Bluetooth paging unit.

illustrated in Fig. 5. For the SCO link and ACL link, different packet types have been defined.

Connection Establishment

A critical design issue in ad hoc radio systems is connection establishment. How do units find each other, and how do they make connections? In Bluetooth, three elements have been defined to support connection establishment: scan, page, and inquiry. A unit in idle mode wants to sleep most of the time to save power. However, in order to allow connections to be made, the unit frequently has to listen whether other units want to connect. In truly ad hoc systems, there is no common control channel a unit can lock to in order to listen for page messages, as is common in conventional (cellular) radio systems. In Bluetooth, a unit periodically wakes up to listen for its identity. However, the explicit identity is not used, but the access code derived from this identity. When a Bluetooth unit wakes up to scan, it opens its sliding correlator which is matched to the access code derived from its own identity. The scan window is a little longer than 10 ms. Every time the unit wakes up, it scans at a different hop carrier. This is required by the regulations, which do not permit a fixed wake-up frequency, and also provides the necessary interference immunity. The Bluetooth wake-up hop sequence is only 32 hops in length and is cyclic. All 32 hops in the wake-up sequence are unique, and they span at least 64 MHz of the 80 MHz available. The sequence is pseudo-random and unique for each Bluetooth device. The sequence is derived from the unit identity. The phase in the sequence is determined by the free-running native clock in the unit. Thus, during idle mode, the native clock is used to schedule wake-up operations. It will be understood that a trade-off has to be made between idle mode power consumption and response time: increasing the sleep time will reduce power consumption, but will prolong the time before an access can be made. The unit that wants to


connect has to solve the frequency-time uncertainty: it does not know when the idle unit will wake up and on which frequency. The burden of solving this uncertainty is deliberately placed at the paging unit because this will require power consumption. Since a radio unit will be in idle mode most of the time, the paging unit should take the power burden. We first assume that the paging unit knows the identity of the unit to which it wants to connect. Then it knows the wake-up sequence and can also generate the access code which serves as the page message. The paging unit then transmits the access code repeatedly at different frequencies: every 1.25 ms; the paging unit transmits two access codes and listens twice for a response (Fig. 6).

Consecutive access codes are transmitted on different hops selected from the wake-up sequence. In a 10 ms period 16 different hop carriers are visited, which represent half of the wake-up sequence. The paging unit transmits the access code on these 16 frequencies cyclically for the duration of the sleep period of the idle unit. If the idle unit wakes up in any of these 16 frequencies, it will

receive the access code and a connection setup procedure follows. However, since the paging unit does not know the phase the idle unit is using, the idle unit can equally well wake up in any of the 16 remaining frequencies in the 32-hop wake-up sequence. Therefore, if the paging unit does not receive a response from the idle unit after a time corresponding to the sleep time, it will transmit the access code repeatedly on the hop carriers in the remaining half of the sequence.³ The maximum access delay therefore amounts to twice the sleep time. When the idle unit receives the page message, it notifies the paging unit by returning a message, which again is the access code derived from the idle unit's identity. Thereafter the paging unit transmits an FHS packet which contains all of the pager's information (e.g., identity and clock). This information is then used by both the paging unit and the idle unit to establish a piconet; that is, the paging unit becomes the master using its identity and clock to define the FH channel, and the idle unit becomes the slave.

The above-described paging process assumes that the paging unit has no knowledge at all of the clock in the idle unit. However, if the units have met before, the paging unit will have an estimate of the clock in the idle unit. When units connect, they exchange their clock information, and the time offsets between their free-running native clocks are stored. This offset is only accurate during the connection; when the connection is released, the offset information becomes less reliable due to clock drifts. The reliability of the offsets is inversely proportional to the time elapsed since the last connection. However, the paging unit can exploit the offset infor-

³ In determining the hop carriers of the second half of the sequence, the paging unit takes into account that the clock in the idle unit also progresses. The remaining half will therefore have one carrier in common with the first half.



■ **Figure 7.** The basic concept of hop selection in Bluetooth.

mation to estimate the phase of the idle unit. Suppose the clock estimate of the idle unit in the paging unit is K . If $f(m)$ is the hop in the wake-up sequence at time m , the paging unit will assume that the idle unit will wake up in $f(K)$. But since in 10 ms it can cover 16 different frequencies, it will also transmit the access code a hop frequencies before and after $f(K)$ or $f(K-8), f(K-7), \dots, f(K), f(K+1), \dots, f(K+7)$. As a result, the phase estimate in the paging unit can be off by -8 or $+7$ while it still covers the wake-up frequency of the unit in idle mode. With a free-running clock accuracy of ± 250 ppm, the clock estimate K is still useful at least 5 hr after the last connection. In this case, the average response time is reduced to half the sleep time.

To establish a connection, the identity of the recipient is required to determine the page message and wake-up sequence. If this information is not known, a unit that desires to make a connection may broadcast an inquiry message that induces recipients to return their address and clock information. With the inquiry procedure, the inquirer can determine which units are in range and what their characteristics are. The inquiry message is again an access code, but derived from a reserved identity (the inquiry address). Idle units also listen to the inquiry message according to a 32-hop inquiry sequence. Units that receive the inquiry message return an FHS packet which includes, among other things, their identity and clock information. For the return of the FHS packet a random backoff mechanism is used to prevent multiple recipients transmitting simultaneously.

During the page and inquiry procedures, 32 hop carriers are used. For pure hopping systems, at least 75 hop carriers must be used. However, during the page and inquiry procedures, only an access code is used for signaling. This access code is used as a direct-sequence code. The processing gain obtained from this direct-sequence code combined with the processing gain obtained from the 32-hop sequence provides sufficient processing gain to satisfy the regulations for hybrid DS/FH systems. Thus, during the page and inquiry procedures the Bluetooth system acts like a hybrid DS/FH system, whereas during the connection it acts as a pure FH system.

Hop Selection Mechanism

Bluetooth applies a special hop selection mechanism. The hop selection mechanism can be considered a black box with an identity and clock in, and a hop carrier out (Fig. 7). The mechanism satisfies the following requirements:

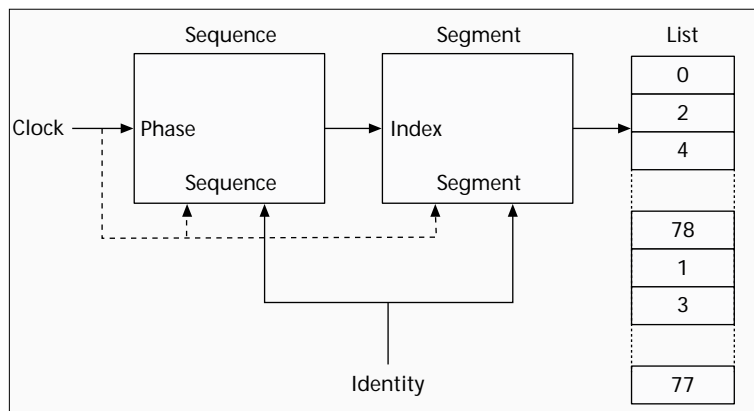
- The sequence is selected by the unit identity, the phase by the unit clock.
- The sequence cycle covers about 23 hours.
- 32 consecutive hops span about 64 MHz of spectrum.
- On average, all frequencies are visited with equal probability.

⁴ For 23-hop systems, a corresponding scheme is constructed with 16-hop segments and a 23-hop list.


- The number of hop sequences is very large.
- By changing the clock and/or identity, the selected hop changes instantaneously.

Note that no extra effort has been taken to make the sequences orthogonal. With only 79 hop carriers, the number of orthogonal sequences is rather limited. The first requirement supports the piconet concept where the master unit defines the hop channel by its identity and clock. The second requirement prevents repetitions in the interference pattern when several piconets are collocated. Repetitive interference is detrimental for synchronous services such as voice. The spanning requirement provides maximal interference immunity by spreading as much as possible over a short time interval. Again, this is most important for voice services. It also provides the desired features for the wake-up and inquiry sequences which are 32 hops in length. Over a larger interval, regulations require that all carriers are visited with equal probability. Since many piconets can coexist in the same area, many hop patterns must be available. This excludes the use of prestored sequences: the sequences are generated on the fly by logic circuitry. Finally, the last requirement provides flexibility to run backward and forward in the sequence by running the clock backward or forward, which is attractive in the page and inquiry procedures. In addition, it supports jumping between piconets: a unit can jump from one piconet to another by merely changing the master parameters (i.e., identity and clock). The latter requirement excludes the use of a memory in the algorithm: only combinatorial logic circuitry is used.

The selection mechanism is illustrated in Fig. 8.⁴ In the first block, the identity selects a 32-hop subsequence with pseudo-random properties. The least significant part of the clock hops through this sequence according to the slot rate (1600 slots/s). The first block thus provides an index in a 32-hop segment. The segments are mapped on the 79-hop carrier list. The carrier list is constructed in such a fashion that even-numbered hops are listed in the first half of the list, odd-numbered hops in the second half of the list. An arbitrary segment of 32 consecutive list elements spans about 64 MHz. For the paging and inquiry procedures, the mapping of the 32-hop segment on the carrier list is fixed. When the clock runs, the same 32-hop sequence and 32 hop carriers will be used. However, different identities will map to different segments and different sequences, so the wake-up hop sequences of different units are well randomized. During the connection, the more significant part of the clock affects both sequence selection and segment mapping: after 32 hops (one segment) the sequence is altered, and the segment is shifted in the forward direction by half its size (16 hops). Segments, each 32 hops in



■ **Figure 8.** The hop selection mechanism; the dashed line for the more significant clock part is used in connection mode only.



■ Figure 9. An example of retransmission operation in Bluetooth.

length, are concatenated, and the random selection of the index changes for each new segment; the segments slide through the carrier list, and on average all carriers are visited with equal probability. Changing the clock and/or identity will directly change the sequence and segment mapping.

Error Correction

Bluetooth includes both FEC and packet retransmission schemes. For FEC, a 1/3-rate code and a 2/3-rate FEC code are supported. The 1/3-rate code merely uses a 3-bit repeat coding with majority decision at the recipient. With the repeat coding, extra gain is obtained due to the reduction of instantaneous bandwidth. As a result, intersymbol interference (ISI) introduced by the receive filtering is decreased. The 1/3-rate code is used for the packet header, and can additionally be applied on the payload of the synchronous packets on the SCO link. For the 2/3-rate FEC code, a shortened Hamming code is used. Error trapping can be applied for decoding. This code can be applied on both the payload of the synchronous packets on the SCO link and the payload of the asynchronous packets on the ACL link. The applied FEC codes are very simple and fast in encoding and decoding operations, which is a requirement because of the limited processing time between RX and TX. This will be further apparent in the next paragraph.

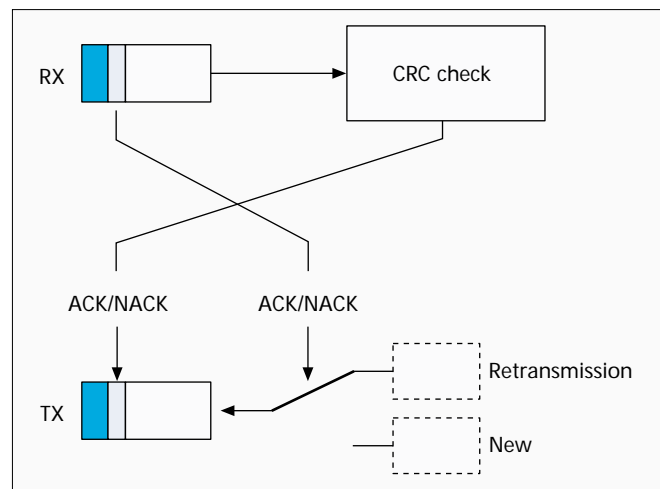
On the ACL link, an ARQ scheme can be applied. In this scheme, a packet retransmission is carried out if the reception of the packet is not acknowledged. Each payload contains a CRC to check for errors. Several ARQ schemes have been considered like stop-and-wait ARQ, go-back- N ARQ, and selective-repeat ARQ [9]. Also, hybrid schemes have been analyzed. However, to minimize complexity, overhead, and wasteful retransmissions, Bluetooth has implemented a fast-ARQ scheme where the sender is notified of the packet reception in the RX slot directly following the TX slot in which the packet was sent (Fig. 9). If the 2/3-rate FEC code is added, a type I hybrid ARQ scheme results. The ACK/NACK information is piggybacked in the packet header of the return packet. There is only the RX/TX switching time for the recipient to determine the correctness of the received packet and creating the ACK/NACK field in the header of the return packet. In addition, the ACK/NACK field in the header of the packet received indicates whether the previously sent payload was correctly received, and thus determines whether a retransmission is required or the next packet can be sent. This process is illustrated in Fig. 10. Due to the short processing time, decoding is preferably carried out on the fly while the packet is received. In addition, the simplicity of the FEC coding schemes speed up the processing. The fast-ARQ scheme is similar to the stop-and-wait ARQ scheme, but the delay has been minimized; in fact, there is no additional delay caused by the ARQ scheme. The scheme is more efficient than go-back- N , since only

failed packets are retransmitted. This is the same efficiency obtained with selective-repeat ARQ, but with reduced overhead: only a 1-bit sequencing number suffices in the fast-ARQ scheme (in order to filter out packets that are correctly received twice due to an error in the ACK/NACK field).


Power Management

In the Bluetooth design, special attention has been paid to reduction of current consumption. In the idle mode, the unit only scans a little over 10 ms every T s where T can range from 1.28 to 3.84 s. Thus, the duty cycle is well below 1 percent. Additionally, a PARK mode has been defined where the duty cycle can be reduced even more. However, the PARK mode can only be applied after the piconet has been established. The slave can then be parked; that is, it only listens to the channel at a very low duty cycle. The slave only has to listen to the access code and the packet header (126 μ s excluding guard time to account for drift) to resynchronize its clock and decide whether it can return to sleep. Since there is no uncertainty in time and frequency (the parked slave is locked to the master, similar to how cordless and cellular phones are locked to their base stations), a much lower duty cycle is achievable. Another low-power mode during connection is the SNIFF mode, in which the slave does not scan at every master-to-slave slot, but has a larger interval between scans.

In the connection state, current consumption is minimized, and only transmitting when information is available prevents wasteful interference. If no useful information needs to be exchanged, no transmission takes place. If only link control information needs to be transferred (e.g., ACK/NACK), a NULL packet without payload is sent. Since NACK is implicit, a NULL packet with NACK does not have to be sent. In longer periods of silence, the master once in a while needs to send a packet on the channel such that all slaves can resynchronize their clocks and compensate for drift. The accuracy of the clocks and the scan window length applied in the slave determines the period of this resynchronization. During continuous TX/RX operations, a unit starts to scan for the access code at the beginning of the RX slot. If in a certain window this access code is not found, the unit returns to sleep until



■ Figure 10. ARQ mechanisms where received ACK/NAK information decides on retransmission and received payload determines transmitted ACK/NAK information.



■ **Figure 11.** *The Bluetooth authentication procedure.*

the next TX slot (for the master) or RX slot (for the slave). If the access code is received (which means the received signal matches the expected access code), the header is decoded. If the 3-bit slave address does not match the recipient, further reception is stopped. The header indicates what type of packet it is and how long the packet will last; therefore, the non-addressed recipients can determine how long they can sleep.

The nominal transmit power used by most Bluetooth applications for short-range connectivity is 0 dBm. This both restricts current consumption and keeps interference to other systems to a minimum. However, the Bluetooth radio specifications allow TX power up to 20 dBm. Above 0 dBm, closed-loop received signal strength indication (RSSI)-based power control is mandatory. This power control only compensates for propagation losses and slow fading. In the uncoordinated environment where ad hoc systems operate, interference-based power control is to say the least doubtful, especially since different types of systems with different power characteristics share the same band. Since power control cannot be coordinated among different systems, it cannot be prevented that certain systems always try to overpower their contenders, and the strongest transmitter will prevail.

Security

Although Bluetooth is mainly intended for short-range connectivity between personal devices, some basic security elements are included to prevent unauthorized usage and eavesdropping. At connection establishment, an authentication process is carried out to verify the identities of the units involved. The authentication process uses a conventional challenge-response routine illustrated in Fig. 11. The claimant (right) transmits its claimed 48-bit address to the verifier (left). The verifier returns a challenge in the form of a 128-bit random number (AU RAND). The AU RAND, the claimant address, and a 128-bit common secret link key form the inputs to a computational secure hash function *E1* based on SAFER+, which produces a 32-bit signed response (SRES). The SRES produced by the claimant is sent to the verifier, which compares this result with its own SRES. Only if the two calculated SRES numbers are the same will the challenger continue with connection establishment. The authentication can be uni- or bidirectional.

In addition to the 32-bit SRES, the *E1* algorithm produces a 96-bit authenticated cipher offset (ACO). This offset is used in the encryption procedure. To prevent eavesdropping on the link, which is a danger inherent to radio communications even if the intended recipient is only at short range, the payload of each packet is encrypted. Encryption is based on stream-ciphering; the payload bits are modulo-2 added to a binary keystream. The binary keystream is generated by a second


hash function *E0* which is based on linear feedback shift registers (LFSRs). When encryption is enabled, the master sends a random number EN RAND to the slave. Before the transmission of each packet, the LFSR is initialized by a combination of this EN RAND, the master identity, an encryption key, and the slot number. Since the slot number changes for each new packet, the initialization is new for each packet. The encryption key is derived from the secret link key, the EN RAND, and the ACO.

The central element in the security process is the 128-bit link key. This link key is a secret key residing in the Bluetooth hardware and is not accessible by the user. The link key is generated during an initialization phase. Two units that want to authenticate each other and establish secure links in the future have to be associated (i.e., provided with the same secret link key). An initialization phase initiated by the user is required to associate two devices. To authorize initialization, the user has to enter an identical PIN in both devices. For devices without a user interface (e.g., headsets), initialization is only possible during a short time window (e.g., after the user has pressed an initialization key). Once the initialization has been carried out, the 128-bit link keys reside in the devices and can from then on be used for automatic authentication without user interaction. In principle, the link key provides an agreement between two units. Thus, to provide security in N units, $N \times (N - 1)/2$ link keys are required. Bluetooth provides methods to reduce the number of keys in certain applications. If a single unit is used by many users (e.g., a printer shared by several users), a single key is used by all users for secure communications to this single unit. In addition, methods are available to use the same encryption key for all slaves in a single piconet.

Bluetooth provides a limited number of security elements at the lowest level. More advanced security procedures (e.g., public keys, certificates) can be implemented at higher layers.

Interpiconet Communications

The Bluetooth system has been optimized to have tens of piconets operate in the same area without noticeable performance degradation. Multiple piconets in the same area are referred to as a *scatternet*. Due to the fact that Bluetooth uses packet-based communication over slotted links, it is possible to interconnect different piconets. This means that units can participate in different piconets. However, since the radio can tune to a single hop carrier only, at any instant in time a unit can communicate in one piconet only. However, the unit can jump from one piconet to another by adjusting the piconet channel parameters (i.e., the master identity and master clock). A unit can also change role when jumping from one piconet to another. For example, a unit can be the master in one piconet at one instant in time, and be a slave in a different piconet at another instant in time. A unit can also be a slave in different piconets. However, by definition, a unit cannot be the master in different piconets, since the master parameters specify the piconet FH channel. The hop selection mechanism has been designed to allow for interpiconet communications: by changing the identity and clock input to the selection mechanism, instantaneously a new hop for the new piconet is selected. In order to make jumps between different piconets feasible, guard time has to be included in the traffic scheduling to account for the slot misalignment of different piconets. In Bluetooth, a HOLD mode has been introduced to allow a unit to temporarily leave one piconet and visit another



■ Figure 12. *The Bluetooth protocol stack.*

(HOLD can also be used as an additional low-power mode when no new piconet is visited during the leave). Traffic scheduling and routing in a scatternet with interpiconet communications is a challenge and still a subject for future study.

Bluetooth Standardization

In the beginning of 1998, a Bluetooth Special Interest Group (SIG) was formed to further expand and promote the Bluetooth concept and establish an industry standard. The SIG promoters are formed by leading manufacturers of the mobile communication industry, portable computer industry, and chip integration industry: Ericsson, Nokia, IBM, Toshiba, and Intel. Version 1.0 of the specification was published in July 1999. Over 1000 companies have signed as adopters of the technology. The Bluetooth technology is royalty-free. A special certification program, including logos, is under development to guarantee Bluetooth interoperability.

The specified protocol stack of Bluetooth is shown in Fig. 12. This article has dealt mainly with the three lower layers:

- The RF layer, specifying the radio parameters
- The baseband layer, specifying the lower-level operations at the bit and packet levels (FEC operations, encryption, CRC calculations, ARQ protocol)
- The link manager (LM) layer, specifying connection establishment and release, authentication, connection and release of SCO and ACL channels, traffic scheduling, link supervision, and power management tasks

The Logical Link Control and Adaptation Protocol (L2CAP) layer has been introduced to form an interface between standard data transport protocols and the Bluetooth protocol. It handles multiplexing of higher-layer protocols, and segmentation/reassembly of large packets. The data stream crosses the LM layer, where packet scheduling on the ACL channel takes place. The audio stream is directly mapped on an SCO channel and bypasses the LM layer. The LM layer, though, is involved in the establishment of the SCO link. Between the LM layer and the application, control messages are exchanged in order to configure the Bluetooth transceiver for the considered application. Above the L2CAP layer, RFCOMM, transmission convergence sublayer (TCS), and other network protocols (e.g., TCP/IP, PPP, OBEX, Wireless Application Protocol) may reside. RFCOMM and TCS are also specified in Bluetooth and provide serial cable emulation and a cordless telephony protocol, respectively. SDP is a service discovery protocol which enables a Bluetooth unit to find the capabilities of other Bluetooth units in range. It discovers which services are available

and the characteristics of these services. This can involve common services like printing, faxing, and so on, as well as more advanced services like teleconferencing, network bridging and access points, e-commerce facilities, and so on. SDP specifically addresses the Bluetooth environment; it does not specify the methods for accessing the service, for which other (non-Bluetooth) protocols can be used.

In addition to protocols which guarantee that two units speak the same language, profiles are defined. Profiles are associated with applications. The profiles specify which protocol elements are mandatory in certain applications. This concept prevents devices with little memory and processing power implementing the entire Bluetooth stack when they only require a small fraction of it. Simple devices like a headset or mouse can thus be implemented with a strongly reduced protocol stack. Profiles are dynamic in the sense that for new applications, new profiles can be added to the Bluetooth specification.

Conclusions

In this article the Bluetooth radio system is presented. The focus is on its capabilities to provide ad hoc radio connectivity. With the restrictions set by regulations, power consumption, lack of coordination, and interference immunity, a robust radio system has evolved which provides a universal wireless interface to a large range of low-cost, portable devices. The article has also described the motivation of the various design choices.

References

- [1] M. Mouly and M.-B. Pautet, *The GSM System for Mobile Communications*, 1992.
- [2] TIA/EIA/IS-136.2, "800 MHz TDMA Cellular-Radio Interface-Mobile Station-Base Station Compatibility — Traffic Channels and FSK Control Channel," Dec. 1994.
- [3] TIA/EIA IS-95B, "Mobile Station-Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular Systems," 1998.
- [4] IEEE 802.11, "Wireless LAN MAC and Physical Layer Specification," June 1997.
- [5] ETSI RES, "High Performance Radio Local Area Network (HIPERLAN) Type 1, Functional Specifications," ETS 300 652, 1996.
- [6] ETSI BRAN, "HIPERLAN Type 2, Functional Specifications," preliminary.
- [7] ETSI RES, "Digital European Cordless Telecommunications (DECT), Common interface Part 1: Overview," ETS 300 175-1, 1996.
- [8] "Personal Handy Phone Standard (PHS)," CRC STD-28, 1993.
- [9] S. Lin and D. J. Costello, *Error Control Coding*, Prentice-Hall, 1983.

Biography

JAAP C. HAARTSEN (jaap.haartsen@erh.ericsson.se) joined Ericsson Mobile Communications in 1991 and has since worked at sites in RTP, the United States, and Lund, Sweden in the area of wireless technology. In Sweden he worked on the foundations of the Bluetooth radio concept. Currently, he is located in Emmen, the Netherlands, where he is working with the Bluetooth system for both current and future applications. Jaap is chair of the Bluetooth air protocol group. He earned M.Sc. and Ph.D. degrees (both with honors) in electrical engineering from Delft University of Technology, the Netherlands. He holds over 25 patents.

Service Advertisement and Discovery:

Enabling Universal Device Cooperation

GOLDEN G. RICHARD III
University of New Orleans

Service advertisement and discovery technologies enable device cooperation and reduce configuration hassles, a necessity in today's increasingly mobile computing environments. This article surveys five competing but similar "service discovery suites" and looks at efforts to bridge the technologies.

Computer users increasingly face the management of many computing devices. One reason is the expansion of computing environments in the home and office, as printers, scanners, digital cameras, and other peripherals are integrated into networked environments. Another reason is the proliferation of mobile devices such as laptop and palm-sized computers, cellular phones, and pagers. Because these devices trade functionality for suitable form factors and low power consumption, they are necessarily "peripheral-poor" and must therefore establish connections to neighboring devices for storage, faxing, high-speed network access, and printing.

It is easy to become frustrated when dealing with the configuration and interaction of such a multitude of devices. Service discovery technologies were developed to reduce this frustration and to simplify the use of mobile devices in a network by allowing them to be "discovered," configured, and used by other devices with a minimum of manual effort.

This article briefly surveys five of the leading technologies in this area. Table 1 lists the features of each technology. Although most of these "service discovery suites" promise similar functionality—namely, reduced configuration hassles, improved device cooperation, and automated discovery of required services—they come at the problem from different philosophical and technical approaches. Since none of these technologies is a superset of the others and none is mature enough to dominate the market, interoperation among them will require bridging mechanisms. The survey concludes with a review of some developments in this area.

BLUETOOTH: PICONETS FOR WIRELESS DEVICES

Bluetooth is a low-power, short-range, wireless radio system being developed by the Bluetooth Special Interest Group, an industry consortium whose member companies include Ericsson, Nokia, and IBM. The radio has a range of 10 meters and provides up to seven 1-megabit-per-second links to other Bluetooth devices. Bluetooth operates in the 2.4-GHz indus-

trial scientific and medical (ISM) band to maximize international acceptance and employs a frequency-hopping system to minimize interference. The low-level communications are detailed in the Bluetooth specification.¹

Bluetooth has a small form factor; complete systems can be as small as 2-cm square. The technology supports both isochronous and asynchronous services. A simple isochronous application might link a cellular phone and wireless headset, where the headset and base are both Bluetooth devices. More complicated applications include automatic discovery of wireless network connections and automatic synchronization of data between several Bluetooth devices.

Figure 1 shows the Bluetooth protocol stack. At the bottom, the radio and baseband layers provide the short-range, frequency-hopping radio platform. The link manager protocol (LMP) handles data link setup and provides authentication and encryption services. The logical link control and adaptation protocol (L2CAP) supports multiplexed connectionless and connection-oriented communication over the LMP layer. L2CAP is proprietary, but other network protocols, such as IP, can be built on top of it. L2CAP is also used by higher level protocols. For example, Figure 1 shows links to the Hayes-compatible AT (ATtention) protocol, which provides a standard interface for controlling remote cellular phones and modems; RFCOMM, which emulates an RS-232 serial interface; a simple object exchange protocol (OBEX), which enhances Bluetooth's interoperability with IrDA; and Bluetooth's service discovery protocol (SDP).

Groups of up to eight Bluetooth devices can form ad hoc networks called *piconets* to communicate, share services, and synchronize data. In each piconet, a master device coordinates the other Bluetooth devices (including setting the 1,600-hops-per-second frequency-hopping pattern). Individual devices can participate in more than one piconet at a time and can be in one of several states:

- *Standby*—the device is conserving power and waiting to connect to another Bluetooth device.
- *Inquire*—the device is searching for nearby Bluetooth devices.
- *Page*—the device is connecting to another Bluetooth device.

Table 1. Features of the five leading service discovery suites.

Feature	Bluetooth	Jini	Salutation	UPnP	SLP
Service discovery	✓	✓	✓	✓	✓
Service announcement		✓	✓	✓	✓
Service registry		✓	✓		✓
Interoperability	✓	✓	✓		✓
Security	✓	✓			✓

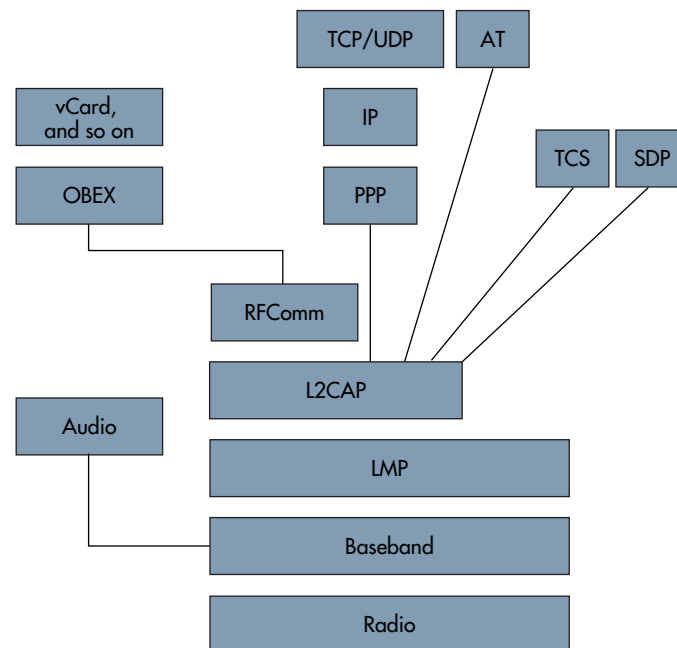


Figure 1. Bluetooth protocol stack. The link manager protocol (LMP) controls link setup and provides encryption and authentication services. The proprietary logical link control and adaptation protocol (L2CAP) provides multiplexed communication over LMP to higher level layers.

- *Connected*—the device is connected to another Bluetooth device.
- *Hold and park*—the device is participating in a piconet with varying degrees of power savings.

The Bluetooth SDP provides a simple API for enumerating the devices in range and browsing available services. It also supports *stop rules* that limit the duration of searches or the number of devices returned. Client applications use the API to search for available services either by service




Figure 2. Jini service discovery entities: clients, lookup servers, and services. In this example, a printer service registers a proxy object with a lookup server, which will serve as a remote control for clients that use the service.

classes, which uniquely identify types of devices (such as printers or storage devices), or by matching attributes (such as a model number or supported protocol). Attributes that describe the services offered by a Bluetooth device are stored as a service record and are maintained by the device's SDP server.

Jini requires each device either to run a Java virtual machine or to associate itself with a device that can execute a JVM on its behalf.

The distinction between service classes and descriptive attributes is not well defined, but service classes generally define broad device categories, such as Printer, ColorPrinter, and PostScriptPrinter, while attributes allow a finer level of description. Manufacturers must eventually standardize these service classes for maximal interoperability between Bluetooth devices.

Unlike higher level service discovery technologies such as Jini, Bluetooth's SDP does not provide a mechanism for using discovered services—specific actions required to use a service must be provided by a higher level protocol. However, it does define a standard attribute ProtocolDescriptorList, which enumerates appropriate protocols for communicating with a service.

Bluetooth devices provide data security through unique 48-bit identifiers, 128-bit authentication keys, and 8- to 128-bit encryption keys. Strong authentication is possible because no international restrictions prevent it, but Bluetooth devices

must negotiate encryption strength to comply with laws restricting encryption. Note that Bluetooth devices must be paired to provide them with matching secret keys that will support authentication. Once paired, Bluetooth devices can authenticate each other and protect sensitive data from snooping. Regardless of encryption strength, Bluetooth's fast frequency-hopping scheme makes snooping difficult.

JINI: MOBILE JAVA CODE

Jini is a service discovery and advertisement system that relies on mobile code and leverages the platform independence of the Java language.² The current Jini implementation is based on TCP and UDP, but implementations based on other network protocols are certainly possible. The major requirements are reliable, stream-oriented communication and a multicast facility. Jini's language-centric approach allows a flexible definition of service; for example, a service can be implemented entirely in software and, after discovery, can be downloaded and executed entirely on the client. Examples of such algorithmic services might include an implementation of a proprietary algorithm for shading a polygon or formatting a document to meet an organizational standard. On the other hand, Jini also requires each device either to run a Java virtual machine or to associate itself with a device that can execute a JVM on its behalf. For example, a Jini "device chassis" might Jini-enable a number of "dumb" devices, making their services available to Jini clients.

Jini entities consist of *services*, *lookup servers* that catalog available services, and *clients* that require services. A service can also be a client; for example, a telescope might provide pictures to a PDA as a service and look for printing services as a client. All service advertisements and requests go through a lookup server. Figure 2 illustrates the discovery and registration process for Jini clients and services.

To register service availability or to discover services, a service or client must first locate one or more lookup servers by using a *multicast request protocol*. This request protocol terminates with the invocation of a *unicast discovery protocol*, which clients and services use to communicate with a specific lookup server. The unicast protocol culmi-

nates in the transfer of an instance of the `ServiceRegistrar` class, a “remote control” for the lookup server. A lookup server can use the *multi-cast announcement protocol* to announce its presence on the network. When a lookup server invokes this protocol, clients and services that have registered interest in receiving announcements of new lookup services are notified.

These three protocols are encapsulated in a set of Jini classes. For example, to find lookup services, a client or service need only create an instance of `LookupDiscovery`.

Jini uses Java’s remote method invocation (RMI) facility for all interactions between either a client or a service and the lookup server (after the initial discovery of the lookup server). Once a lookup server has been discovered and an instance of `ServiceRegistrar` is available, services can register their availability, and clients can search for needed services by invoking `ServiceRegistrar` methods.

Jini associates a proxy, or *remote control object*, with each service instance. A service advertises its availability by registering its object in one or more lookup servers via the `register()` method. This method takes several arguments, including an instance of `ServiceItem`, which contains a universally unique identifier for the service, its attribute set, and its remote control object. This object may either implement the service entirely (in the case of an algorithmic service such as the implementation of a polygon-shading algorithm), or provide methods for accessing the service over the network. The `leaseduration` parameter of `register()` specifies the service’s intended lifetime. The service is responsible for renewing the lease within the time specified to maintain its listing. The lookup server is free to adjust the lease time, which is returned in a `ServiceRegistration` object.

When a service first contacts a lookup server, the server generates a unique identifier for it; the service uses this ID in all future registrations. The service identifier lets clients request a specific service explicitly and recognize when services reported by different lookup servers are identical.

To use a service, a device must first secure an instance of the proxy object for it. From a client point of view, the location of the service proxied by this remote control object is unimportant, because the object encapsulates the location of the service and the protocol necessary to operate it.

Clients use the `lookup()` method in `ServiceRegistrar` to discover services. This method takes a single argument, an instance of `ServiceTemplate`. The




Figure 3. Salutation architecture. Salutation managers are service brokers, isolated by transport managers from the details of specific network transport protocols.

`ServiceTemplate` constructor takes several arguments. The first is the service identifier. If the service identifier is null, then arrays of types (Java classes, typically interfaces) and attributes (attribute objects) are used to match services. A service matches if its class matches one of the classes in the types array and if, for each of the attribute objects, all non-null members match one of the service’s registered attributes. The return value from `lookup()` is an instance of `ServiceMatches`, which contains an array of remote control objects for the services that match. Finally, the `notify()` method allows a client to request an asynchronous notification when services matching a `ServiceTemplate` instance become available. This method uses Jini’s distributed events mechanism, which extends Java’s infrastructure for eventing across JVMs.

Jini depends on Java’s security model, which provides tools like digital certificates, encryption, and control over mobile code activities such as opening and accepting socket connections, reading and writing to specific files, and using native methods. Systems administrators can establish different policies depending on where the Java code originated (for example, the local file system or a remote machine).

SALUTATION: A NETWORK-INDEPENDENT ARCHITECTURE

Salutation is an architecture for service discovery under development by the Salutation Consortium, which includes members from both industry and academia.³ The consortium’s goal is to build a royalty-free architecture for service advertisement and discovery that is independent of a particular network transport.

Figure 3 shows the three fundamental components in the Salutation architecture: *functional units*,

salutation managers, and *transport managers*. From a client's point of view, a functional unit defines a service. Functional units already specified or under consideration by the Salutation Consortium include printing, faxing, and document storage. There is also work on a functional unit specification to allow discovery of Hewlett-Packard JetSend-enabled devices. The specifications define attributes that

Salutation requires a network transport protocol that supports reliable, stream-oriented communication.

characterize a service (for example, in the case of a printer, double-sided capability, color, and so on).

The functional unit Doc Storage defines file attributes that can be used to find information in temporary or long-term storage. For example, a client can search for operating system-specific drivers or software necessary to interact with a newly discovered device. The client simply queries a Salutation manager for the necessary Doc Storage functional unit, extracts the application or device driver, and installs it, thus providing limited code mobility.

Salutation managers function as service brokers; they help clients find needed services and let services register their availability. Services can register and unregister functional units with the local Salutation manager by using the API calls `slmRegisterCapabilities()` and `slmUnregisterCapabilities()`, respectively. A client can use the `slmSearchCapability()` call to determine if Salutation managers have registered specific functional units. Under the current version of the architecture, applications can query only the local Salutation manager. Future versions will allow remote Salutation managers to be specified. Once a functional unit is discovered, `slmQueryCapability()` can be used to verify that a functional unit has certain capabilities. The API also includes calls for initialization/version checking, availability checking, and communication between clients and services. (An API simulator is available at <http://www.salutation.org/simulate.htm>.)

Salutation managers fill a role similar to lookup servers in Jini, but they can also manage the connections between clients and services. A Salutation manager can operate in one of three "personalities":

- In *native* personality, Salutation managers are used only for discovery. They establish a connection between a client and service but perform no further operations on the data stream.
- The *emulated* personality is similar to the native personality in that Salutation managers set up the connection, but in this case they transfer native data packets encapsulated in Salutation manager protocol format, providing a bridge when no common message protocol exists between client and service. The Salutation manager is ignorant of the semantic content of the data stream between client and service.
- In *Salutation* personality, Salutation managers establish the connection between client and service, and they also mandate the specific format of the data transferred. The Salutation architecture defines the data formats.

A transport manager isolates the implementation of the Salutation manager from particular transport-layer protocols and thereby gives Salutation network transport independence. To support a new network transport requires a new transport manager to be written, but does not require modifications to the Salutation manager. Like Jini (and UPnP), Salutation requires a network transport protocol that supports reliable, stream-oriented communication. Initial implementations are based on IP and IrDA because of their widespread use.

Transport managers also locate the Salutation managers on their respective network segments via either multicast, static configuration, or reference to a centralized directory. Discovery of other Salutation managers allows a particular Salutation manager to determine which functional units have been registered and to allow clients access to these remote services. Communication between Salutation managers is based on remote procedure call (RPC). This interaction between remote Salutation managers contrasts with other registry-based service discovery mechanisms (for example, Jini and Service Location Protocol), where clients would be responsible for locating remote registries.

The Salutation specification currently does not address security issues.

A lightweight version of Salutation, called Salutation-Lite, has been developed for resource-limited devices. It is based primarily on IrDA to leverage the large number of infrared-capable devices. Salutation-Lite focuses primarily on service discovery. It uses the functional units OpEnvironment and Display to describe the operating system, processor

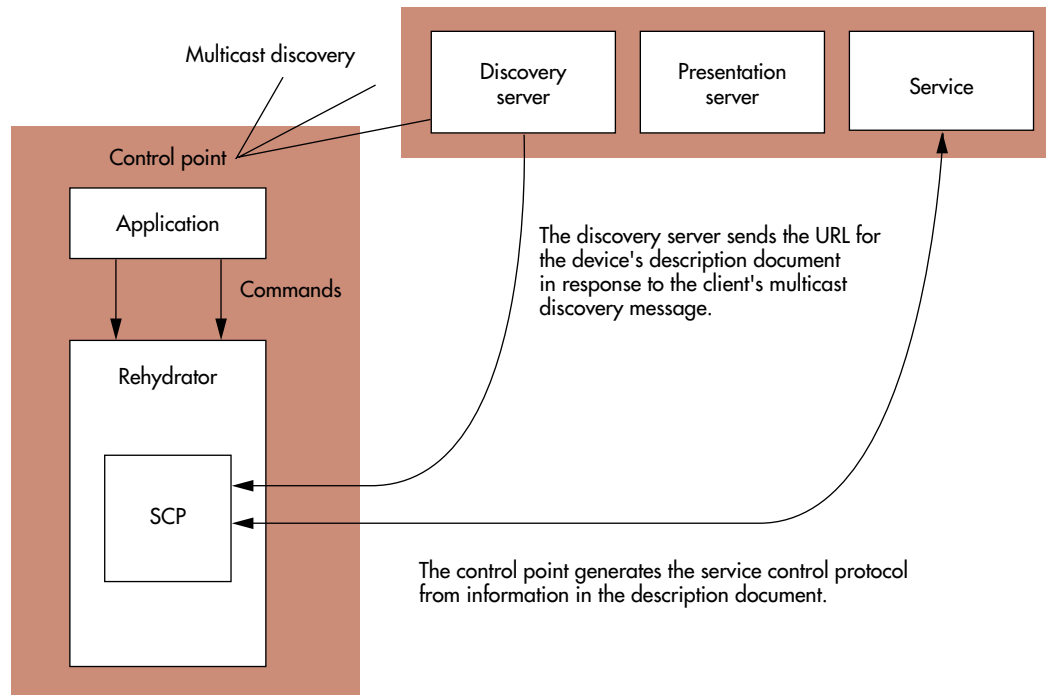


Figure 4. Interaction between a client (control point) and a service in UPnP. The control point discovers the device by sending a multicast message. The device responds with a URL pointing to its description document, which the control point can download for pertinent information, including a URL to which control messages can be sent and the protocol for interacting with the device through this control URL. The “rehydrator” converts generic commands into device-specific control messages.

class, amount of memory, and display characteristics of palm-sized devices. By noting the particular characteristics of the device, servers can provide appropriate drivers and software wirelessly.

Salutation-Lite implementations can be downloaded free from the Salutation website at <http://www.salutation.org>.

UPnP: XML FOR A WEB-BASED ARCHITECTURE

UPnP is a proposed architecture for service advertisement and discovery supported by the UPnP Forum, headed by Microsoft. Unlike Jini, which depends on mobile code, UPnP aims to standardize the protocols used by devices to communicate, using XML. The UPnP specification⁴ is still in a preliminary stage; major issues like security have not yet been addressed.

UPnP’s device model is hierarchical. In a compound device (for example, a VCR/TV combo), the *root device* is discoverable, and a client (called a *control point*) can address the individual subdevices (for example, a tuner) independently. Virtual Web servers

in the device act as entry points for interacting with and controlling it. Devices that don’t speak UPnP directly are called *bridged devices*. They can be integrated into a UPnP network in a manner similar to the integration in a Jini device chassis: A bridge maps between UPnP and device-native protocols.

The UPnP specification describes device addressing, service advertisement and discovery, device control, eventing, and presentation. The eventing facility allows clients to watch for significant changes in the state of a discovered service. It functions similarly to Jini’s distributed event facility. Presentation allows a client to obtain a GUI for a discovered device through one of the device’s virtual Web servers. Several protocols support these functions:

- AutoIP,⁵ a simple protocol that allows devices to dynamically claim IP addresses in the absence of a DHCP server;
- Simple service discovery protocol (SSDP), the UPnP mechanism for service discovery and advertisement;




Figure 5. SLP entities: user agents, directory agents, and service agents. UAs discover services on behalf of applications, either via a DA or directly through an SA. In this example, a laptop and desktop are clients seeking services. A plotter and LCD projection system are services advertising their availability.

- Simple object access protocol (SOAP),⁶ a protocol for remote procedure calls based on XML and HTTP that is used for device control after discovery; and
- Generic Event Notification Architecture (GENA), a UPnP subscription-based event notification service based on HTTP.

When devices are introduced into a network, they multicast “alive” messages to control points. When they wish to cancel availability of their services, they send “byebye” messages. In SSDP, each service has three associated IDs—service type, service name, and location—which are multicast when services are advertised. Any of these IDs can also be used to search for services.

To search, a control point sends a UDP multicast request to the network, as shown in Figure 4. Matching services send unicast responses to the client. These responses contain URLs, each pointing to an XML *description document* that describes a service. A description document contains several important items:

- A *presentation URL* allows entry to a device’s root page, which provides a GUI for device control.

- A *control URL* is the entry point to the device’s control server, which accepts device-specific commands to control the device.
- An *event subscription URL* can be used by clients to subscribe to the device’s event service. The client provides an *event sink URL* in the subscription request. Significant state changes in the device result in a notification to the client’s event sink URL.
- A *service control protocol definition* describes the protocol for interacting with the device.

The service control protocol (SCP) definition allows APIs to be converted to device-specific commands, shielding the application level from details of particular devices. After retrieving the description document, a UPnP component on the control point called the *rehydrator* is “plumbed” with a definition of the device’s SCP. This component then sends device-specific commands via the device’s control URL. SOAP is used for this interaction.

SSDP is similar to the Internet Engineering Task Force’s service location protocol, but it lacks a query facility that can search for services by attributes. Further, SLP incorporates security measures and can interact with the IETF standards-track dynamic host configuration protocol (DHCP)⁷ and the lightweight directory protocol (LDAP).⁸ Finally, SSDP specifications currently limit discovery to a single subnet. Since UPnP does not use a registry, it is also likely to generate significantly more network traffic than SLP.

SLP: A PROPOSED IETF STANDARD

Service location protocol is an IETF protocol for service discovery and advertisement.⁹ It is currently at the “proposed standard” stage along the IETF standards track. Unlike Jini, Salutation, and UPnP, which all aspire to some degree of transport-level independence, SLP is designed solely for IP-based networks. It provides a set of C and Java bindings that provide service discovery and advertisement functions to application software.

SLP comprises three entities: *service agents* (SAs), *user agents* (UAs), and *directory agents* (DAs). SAs advertise the location and attributes of available services, while UAs discover the location and attributes of services needed by client software. UAs can discover services by issuing a directory-like query to the network. DAs cache information about available services. Unlike Jini, SLP can operate without directory servers. The presence of one or more DAs can substantially improve perfor-

mance, however, by reducing the number of multicast messages and the amount of network bandwidth used. In fact, if DHCP is used to configure SLP agents with the location of DAs, then multicast is completely unnecessary. SLP also interoperates with LDAP, so services registered with an SLP DA can be automatically registered in an LDAP directory. This eliminates the need to reconfigure clients that already discover services using LDAP.

SLP has several mechanisms for discovering DAs:

- In passive discovery, SAs and UAs listen for multicast announcements from DAs, which periodically repeat these advertisements.
- In active discovery, SAs and UAs multicast SLP requests or use DHCP to discover DAs. When a DA is present, SAs and UAs use unicast communication to, respectively, register their services and find appropriate services.

In the absence of DAs, UAs multicast requests for service and receive unicast responses directly from the SAs that control matching services. This tends to increase bandwidth consumption, but provides a simpler model, appropriate for small networks (such as a home LAN).

SLP services are advertised through a service URL, which contains all information necessary to contact a service. Clients use the service URL to connect to the service. The protocol used between the client and server is outside the scope of the SLP specification. This separation is similar to Bluetooth, where the SDP does not specifically address how devices will communicate.

Service templates define an attribute set for each service type (a printer, for example).¹⁰ The attributes include a specification of the attribute types and information about default and allowed values; they are used to differentiate between services of the same type and to communicate configuration information to UAs.

SLP doesn't define the protocols for communication between clients and services, and so its security model concentrates on preventing the malicious propagation of false information about service locations. SAs can include digital signatures when registering so DAs and UAs can verify their identity. Digital signatures can also be required when DAs advertise their availability, allowing UAs and SAs to avoid rogue DAs (that is, those without a proper signature). As with Jini, setting up the security features of SLP requires some configura-

tion effort, but the effort can be well worth it, particularly in open environments.

BRIDGING THE TECHNOLOGIES

For service discovery to become pervasive, either a single service discovery technology must dominate or the most commonly used technologies must be made interoperable. Currently, bridging seems to be the most promising prospect for interoperability.

Implementations of certain low-level functions of service discovery (such as discovering registries) are interchangeable.

Implementations of certain low-level functions of service discovery (such as discovering registries) are interchangeable. For example, the Salutation Consortium uses SLP for service discovery beyond the local subnet. This lets the Salutation Manager search for SLP DAs, and then use SLP to register functional units and search for requested services.

A Jini-SLP bridge has also been developed, which allows services lacking a JVM to participate in Jini systems.¹¹ The heart of the Jini-SLP bridge is a special SLP UA that registers the availability of "Jini-capable" SLP SAs. To do this, Jini-capable SLP services advertise the availability of a Jini driver factory. The UA discovers all SAs with driver factories and registers them with one or more Jini lookup services. When a Jini client needs one of the registered SAs, it downloads the driver factory from the lookup server and uses it to instantiate a Java object to drive the service. Note that the SLP SAs are *not* required to host a Java virtual machine—the Java code installed on the SAs is static. Similar schemes are possible for the other technologies; for example, it should be possible to Jini-enable UPnP services in this way.

Miller and Pascoe¹² describe mapping Salutation to Bluetooth SDP to take advantage of Bluetooth's wireless capability. Two approaches are considered: The first maps the Salutation APIs to Bluetooth SDP by implementing Salutation on top of Bluetooth; the second uses a Bluetooth transport manager and essentially replaces Bluetooth SDP with Salutation. This approach will also

work with other schemes, like Jini. Bluetooth is a particularly attractive target for interoperability, primarily because of its wireless capability. Because of this, additional interoperability efforts between Bluetooth and other service discovery technologies seem inevitable.

Each service discovery technology has advantages and disadvantages. Currently, interoperability efforts are perhaps the most important force in service discovery, since it is very unlikely that device manufacturers will embrace multiple service discovery technologies on low-cost devices. ■

ACKNOWLEDGMENTS

Thanks to Sumi Helal of the University of Florida for sparking my interest in this area. Much of the material in this article is derived from a tutorial he invited me to create for IPCCC 2000 in Phoenix. Many thanks to Erik Guttman of Sun Microsystems for clarifying the differences between SLP and SSDP and going way beyond the call of duty in critiquing early versions of this article. David La Motta and Kirk Perilloux were kind enough to read early versions and offer suggestions. Finally, my personal editor (and mate) Christine Ciarmello-Richard was gracious enough to lend her critical eye, as always.

REFERENCES

1. *Specification of the Bluetooth System*; available at <http://www.bluetooth.com/developer/specification/specification.asp>.
2. K. Arnold et al., *The Jini Specification*, Addison-Wesley Longman, Reading, Mass., 1999.
3. *Salutation Architecture Specification*; available online at <http://www.salutation.org/specordr.htm>.
4. *Universal Plug and Play specification v1.0*; available online at <http://www.upnp.org/>.
5. R. Troll, "Automatically Choosing an IP Address in an Ad-Hoc IPv4 Network," IETF Internet draft, work in progress, Mar. 2000.
6. Simple Object Access Protocol (SOAP) 1.1, W3C Note; available online at <http://www.w3.org/TR/SOAP>.
7. R. Droms, "Dynamic Host Configuration Protocol," IETF RFC 2131, Mar. 1997; available online at <http://www.dhcp.org/rfc2131.html>.
8. M. Wahl, T. Howes, and S. Kille, "Lightweight Directory Access Protocol, version 3," IETF RFC 2251, Dec. 1997; available online at <http://www.rfc-editor.org/rfc/rfc2251.txt>.
9. E. Guttman, "Service Location Protocol: Automatic Discovery of IP Network Services," *IEEE Internet Computing*, vol. 3, no. 4, July/Aug. 1999, pp. 71-80.
10. E. Guttman, C. Perkins, and J. Kempf, "Service Templates and Service: Schemes," IETF RFC 2609, June 1999; available online at <http://www.rfc-editor.org/rfc/rfc2609.txt>.
11. E. Guttman and J. Kempf, "Automatic Discovery of Thin Servers: SLP, Jini and the SLP-Jini Bridge," *Proc. 25th Ann. Conf. IEEE Industrial Electronics Soc. (IECON 99)*, IEEE Press, Piscataway, N.J., 1999.
12. B. Miller and R. Pascoe, "Mapping Salutation Architecture APIs to the Bluetooth Service Discovery Layer," white paper; available online at <http://www.salutation.org/whitepaper/BtoothMapping.pdf>.

Golden G. Richard III is an assistant professor of computer science at the University of New Orleans in Louisiana. His research interests include mobile computing, wireless networking, operating systems, and fault tolerance. He is on the executive committee of the IEEE Technical Committee on the Internet, a member of the IEEE and the ACM, and liaison to the University of New Orleans for Usenix's Educational Outreach Program.

Readers may contact the author at golden@cs.uno.edu.

How to Reach IC

Writers

We welcome submissions about Internet application technologies. For detailed instructions and information on peer review, *IEEE Internet Computing's* author guidelines are available online at <http://computer.org/internet/edguide.htm>.

Letters to the Editor

Please send letters via e-mail to internet-computing@computer.org.

Reuse Permission

For permission to reprint an article published in *IC*, contact William J. Hagen, IEEE Copyrights and Trademarks Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855-1331; w.hagen@ieee.org. Complete information is available at <http://computer.org/permission.htm>. To purchase reprints, visit <http://computer.org/author/reprint.htm>.

APPENDIX F

**Mc
Graw
Hill**

**[Thoroughly Explains More Than
1,400 Networking Concepts]**


Encyclopedia of Networking & Telecommunications



**CD contains more than
5,000 hyperlinks, and
a complete set of
cross-referenced
Internet engineering
documents**

**Covers emerging technologies:
all-optical networks,
broadband access,
wireless computing, QoS,
.NET, and more**

**Addresses vendor-specific
technologies: Microsoft®,
Cisco®, IBM®, Juniper®,
Nortel®, Sun®, and others**



Tom Sheldon

Certified Network Engineer and author of the best-selling
Encyclopedia of Networking, Electronic Edition

OSBORNE 
ROKU EXH. 1002



McGraw
Hill

APPENDIX F

Encyclopedia

of Networking & Telecommunications

The Most Complete Volume of Networking Technologies Available

TOPICS COVERED:


ASP (Application Service Provider)
Bluetooth
Cryptography
Distributed Computer Networks
Embedded Systems
Hacking and Hackers
InfiniBand
Java
Linux
Load Balancing
Mobile Computing
NAS (Network Attached Storage)
Network Processors
Optical Networks
Outsourcing
PKI (Public Key Infrastructure)
QoS (Quality of Service)
SAN (Storage Area Network)
Switching Fabrics
Transaction Processing
UNIX
Webcasting
XML



Authoritative and up-to-date, this all-encompassing book and CD-ROM package is filled with thousands of explanations and analyses of core and cutting-edge networking and telecommunications topics—from Abilene to QoS to ZAWS. Extensive cross-referencing throughout helps you understand the relationship among the technologies.

This is a must-have resource for every network professional, as well as technology investors, marketing managers, head hunters, technology writers, and anyone interested in networking.

The book also includes the most comprehensive guide to Internet engineering documents (RFCs) available today. *The McGraw-Hill Encyclopedia of Networking & Telecommunications* reflects the latest in networking and Internet technologies.



On the value-packed CD-ROM

- Complete, fully searchable version of the book with thousands of hyperlinks to related topics in the book
- External hyperlinks to author-selected Web sites for further information
- Illustrations of complex networking topics

Praise for the previous edition:

"In the rapidly converging disciplines of voice and data networks, I haven't found any other source that provides the coverage that [this book] provides."

—Randy Johnson,
Applications Engineer,
Nokia IP Telephony
Business Unit

"The best reference I could find as a student learning about networks."

—Joe Higgins,
Southwest Memorial
Hospital, Director of
Education/Telemedicine

"I interface daily with network engineers. This encyclopedia has helped me 'keep pace' with the engineering units."

—Steve Goldman,
Chief Technical Specialist,
Empire Blue Cross and
Blue Shield

ABOUT THE AUTHOR:

Tom Sheldon is a Certified Network Engineer and has been building networks for nearly three decades. In that time, he has operated a network testing lab and constructed networks for Lockheed Space Operations. He is the author of more than 30 highly-acclaimed technical books, including *Encyclopedia of Networking*, *Electronic Edition*, *Windows NT Security Handbook*, and *Microsoft Internet Information Server*. Tom maintains the Linktionary.com Web site.

OSBORNE

REQUIRED READING for the Information Age

A Division of The McGraw-Hill Companies

\$69.99 USA

£51.99 UK

NETWORKING



7 83254 03096 1

www.osborne.com

Book P/N 0-07-212270-6 of
ISBN 0-07-212005-3



9 780072 120059

90000



McGraw-Hill

Encyclopedia of Networking & Telecommunications

Tom Sheldon

Osborne/McGraw-Hill

New York Chicago San Francisco
Lisbon London Madrid Mexico City
Milan New Delhi San Juan
Seoul Singapore Sydney Toronto

Osborne/McGraw-Hill
2600 Tenth Street
Berkeley, California 94710
U.S.A.

To arrange bulk purchase discounts for sales promotions, premiums, or fund-raisers, please contact Osborne/McGraw-Hill at the above address. For information on translations or book distributors outside the U.S.A., please see the International Contact Information page immediately following the index of this book.

McGraw-Hill Encyclopedia of Networking & Telecommunications

Copyright © 2001 by The McGraw-Hill Companies. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

1234567890 DOC DOC 01987654321

Book p/n 0-07-212005-3 and CD p/n 0-07-212005-3
parts of
ISBN 0-07-212005-3

Publisher

Brandon A. Nordin

Vice President & Associate Publisher

Scott Rogers

Acquisitions Editors

Wendy Rinaldi and Ann Sellers

Project Editor

Lisa Wolters-Broder

Acquisitions Coordinator

Timothy Madrid

Technical Editor

Dan Logan

Copy Editor

Dennis Weaver

Proofreaders

Linda and Paul Medoff

Indexer

Jack Lewis

Computer Designers

Michelle Galicia

Tara A. Davis

Illustrator

Michael Mueller

Series Design

Peter F. Hancik

Cover Design

Amparo del Rio

This book was composed with Corel VENTURA™ Publisher.

Information has been obtained by Osborne/McGraw-Hill from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Osborne/McGraw-Hill, or others, Osborne/McGraw-Hill does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from use of such information.

Service Advertising and Discovery

Services on networks can be advertised so that users can discover them. A number of protocols and schemes are available to support service advertising and discovery. For example, service advertising and discovery is important as mobile devices and mobile wireless devices proliferate on networks. These devices may connect to networks at varying locations. A service discovery and advertising protocol is an important tool to help these devices find services on the network wherever they connect, and to let other network users know about the services they are offering.

Keep in mind that as networks evolve, a variety of services will be offered. For example, network services such as file, print, and applications services can be advertised to “foreign” mobile users who temporarily connect to a network. But other possibilities exist, especially in the wireless realm. For example, an airport could have numerous small wireless networks that are limited in range to about 10 or 20 feet. As you walk into the range of one of these networks, various service advertisements appear on your portable device. These may be advertisements for peripherals like printers that you can use, but they could also be commercial advertisements.

If you are familiar with instant messaging, you are familiar with service advertising protocols. When a person in your “buddy list” signs on to the network, you receive an alert and you can start chatting with them over the network via special chat software, Internet phone, or videoconferencing software. In the case of the airport wireless networks mentioned previously, advertising protocols can alert you to friends who are located in the general vicinity. Their wireless device is advertising their personal ID and your wireless device listens for IDs and looks them up in your personal address book. If a friend is nearby, your device gets excited and starts beeping.

Two earlier advertising services that were developed for LAN environments include SAP (Service Advertising Protocol) and NetBIOS (Network Basic Input/Output System). These are discussed under their own heading.

A number of new approaches have been developed to provide enhanced service advertising and discovery in dynamic network environments such as the wireless and mobile computing networks, where devices frequently connect and disconnect from the network. When a device comes online, it advertises its services or listens for advertisements of available services.

One technique a device may use to locate a service on the network is to send out a multicast packet that contains a service request. Network devices that are providing services listen for multicast packets and then determine whether they can satisfy the request for services being made by the client. If so, the service will respond to the client with a positive message.

Here are some architectures and schemes related to service discovery, advertising, and acquisition:

- **Salutation** The Salutation architecture is a royalty-free service discovery and service management product from the Salutation Consortium, a nonprofit corporation. Salutation is an open standard, independent of operating system, communications protocol, hardware platform, or vendor-imposed limitations. It was created to provide service

discovery for a broad range of network appliances and equipment in a platform-, OS-, and network-independent environment. Devices can use it to advertise and describe their capabilities and discover the capabilities of other devices by using search features.

- **SLP (Service Location Protocol)** SLP is an IETF standard designed to make it easy for network clients to discover the available services on a network and learn information about the configuration of those services. Many vendors support SLP in their operating systems, including Apple, IBM, Novell, and Sun Microsystems. The IETF Service Location Working Group is developing SLP and similar services. See "SLP (Service Location Protocol)."
- **Microsoft.NET** The Microsoft.NET platform for Web Services is a development environment based on building applications with "Web Services." The technique is similar to building distributed objects, but is based on HTTP and XML. Data is represented with XML and delivered in SOAP (Simple Object Access Protocol) messages via HTTP. A language called WSDL (Web Services Description Language) is used to describe services. An XML-based protocol called Disco is used to discover services at a site and a mechanism called UDDI (Universal Description, Discovery, and Integration) defines how to advertise services and how Web Service consumers can find services. See Microsoft.NET.
- **SSDP (Simple Service Discovery Protocol)** SSDP is a Microsoft service location protocol that is part of Microsoft's Universal Plug and Play (UpnP) initiative. It is oriented toward home networks. Like SLP, it enables devices to request information about services on a network and to advertise their presence and the services they offer.
- **Bluetooth** This is a wireless connectivity specification that enables electronic devices to talk spontaneously and allows instant wireless connectivity between computers, mobile phones, and portable devices. Bluetooth includes its own service discovery protocol that locates services offered by devices within the vicinity of a user's Bluetooth device. Currently, Bluetooth's service discovery protocol is being mapped to the Salutation architecture. See "Bluetooth."
- **Jini** This is a Java-based technology defined by Sun Microsystems. When Jini-enabled devices connect to networks, they establish impromptu Java-oriented networks that let users immediately access network resources and services. The technology is designed to support any device that "passes digital information in or out" according to Sun. Devices register with the network when they connect, which makes them available to other devices. For example, when a printer is attached and gets registered, it makes its driver available on the network and this driver gets downloaded to clients when they need to use the printer.
- **JetSend (Hewlett-Packard)** JetSend is code that is embedded in devices to allow them to directly exchange information. Devices become either senders or receivers. JetSend gives devices the intelligence to know their own capabilities and negotiate the best way to exchange information with other devices. No external operating systems need to get involved. No special drivers are needed to connect with other devices. All JetSend devices can immediately communicate. JetSend is a transport-independent protocol that

works across any bidirectional transport, including TCP/IP, IR, IEEE 1394, and others. It is ideal for PDAs, digital cameras, copiers, network-attached printers and scanners, fax machines, and other devices.

- **Inferno by Lucent Technologies** A real-time network operating system that provides a software infrastructure for creating distributed network applications. Inferno is more like a file system that operates over a variety of transport protocols. It is designed to provide connectivity over the Internet, public telephone networks, cable television, and satellite broadcast networks. Inferno includes network and security protocols. It has a very small memory footprint and can be used as a stand-alone OS on information appliances.

A lot of the work being done in this area is for home networking and network appliance configuration. In particular, Jini and Microsoft's UPnP are designed to help devices connect and cooperate.

The IETF Resource Capabilities Discovery (rescap) Working Group is developing services that distribute information about resources or services to the global Internet. The IETF Service Location Protocol (svrloc) Working Group has developed procedures for discovering services.

Related Entries Bluetooth; Directory Services; Embedded Systems; Home Networking; Instant Messaging; Java; Microsoft.NET; Mobile Computing; Network Appliances; Search and Discovery Services; *and* SLP (Service Location Protocol)

Linktionary!—Tom Sheldon's Encyclopedia of Networking updates	http://www.linktionary.com/s/service_advertising.html
Salutation Consortium	http://www.salutation.org
IETF Working Group: Service Location Protocol (svrloc)	http://www.ietf.org/html.charters/svrloc-charter.html
IETF Working Group: Resource Capabilities Discovery (rescap)	http://www.ietf.org/html.charters/rescap-charter.html
Microsoft (search for SSDP)	http://www.microsoft.com/
Sun Microsystems JINI network technology	http://www.sun.com/jini/
The Official Bluetooth Web site	http://www.bluetooth.com/
Jetsend home page	http://www.jetsend.com

Service Providers and Carriers

Anyone with an Internet access account and a telephone is familiar with service providers and carriers. You pay them money every month. However, "service providers" and "carriers" are broad categories. This section describes the different types of service providers and carriers and the service they offer.

In the beginning, at least in the United States, there was one phone company: AT&T. The ILECs (*incumbent local exchange carriers*) are the result of the breakup of AT&T in 1984. That breakup created seven independent RBOCs (Regional Bell Operating Companies). These included Pacific Bell, NYNEX, GTE, and others, but mergers and consolidations have changed

Designed for

Microsoft®
Windows NT®
Windows 95

One-Stop Reference

*The Essential Guide for Administrators,
Systems Engineers, and IS Professionals*

APPENDIX G



Running Microsoft® **Windows NT Server 4.0**

Charlie Russel and
Sharon Crawford

Microsoft® Press

RUNNING

Microsoft® Windows NT® Server 4.0

The Essential Guide
for Administrators,
Systems Engineers,
and IS Professionals

*Charlie Russel and
Sharon Crawford*

APPENDIX G

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 1997 by Charlie Russel and Sharon Crawford

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data pending.

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 QMQM 2 1 0 9 8 7

Distributed to the book trade in Canada by Macmillan of Canada, a division of Canada Publishing Corporation.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office. Or contact Microsoft Press International directly at fax (206) 936-7329.

Macintosh is a registered trademark of Apple Computer, Inc. Intel is a registered trademark of Intel Corporation. Microsoft, Microsoft Press, MS-DOS, Windows, Windows NT, Windows NT Server, and Windows NT Workstation are registered trademarks, and ActiveX, BackOffice, and FrontPage are trademarks of Microsoft Corporation. Java is a trademark of Sun Microsystems, Inc. Other product and company names mentioned herein may be the trademarks of their respective owners.

Acquisitions Editor: David J. Clark
Project Editor: Sigrid Anne Strom
Technical Editor: Jim Fuchs

CHAPTER 8

Disk Configuration

What does a server really *do* on a network? Probably the single most important service to users is storing files and then making them available on demand. File services are *the* area that will get you, the system administrator, into the most hot water if there is a failure. It's your job to make sure these conditions are met for your network:

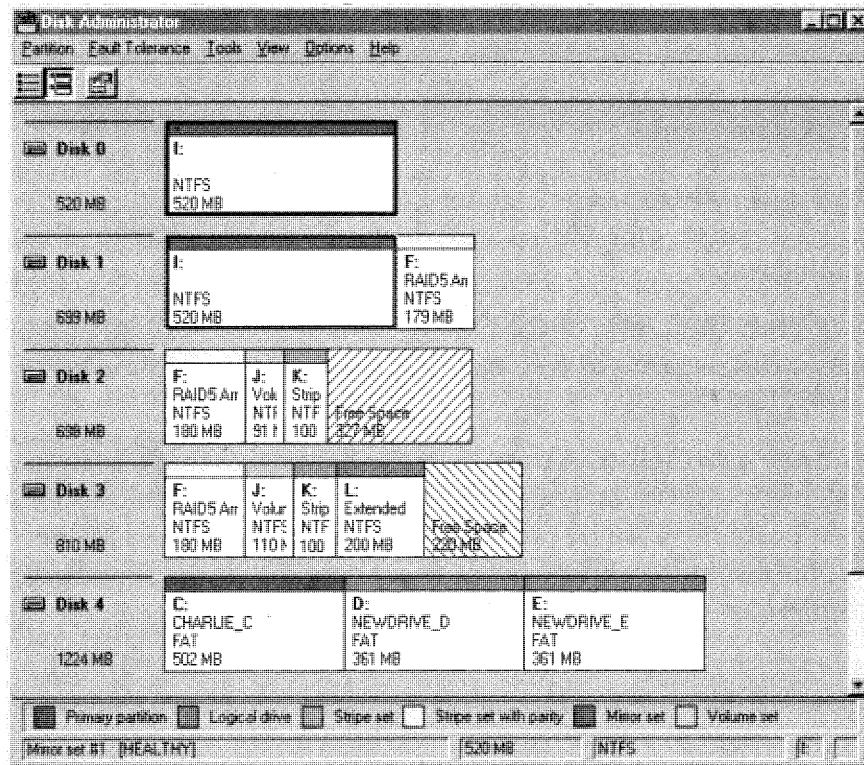
- ◆ There is sufficient hard drive space available.
- ◆ All of the files are backed up.
- ◆ Procedures are in place to recover files in the event of a system failure.
- ◆ Security procedures are sufficient to prevent unauthorized access to files but don't create hardship for normal users.

All of these are important parts of providing quality file services; and, if you've done your job correctly, they'll go completely unnoticed, barring some natural disaster. (If you're looking for glory, you're probably in the wrong job.)

In this chapter, we'll discuss configuring your storage subsystem using the built-in Disk Administrator application of Microsoft Windows NT Server version 4. With Disk Administrator, you can easily manage your disk storage to provide a flexible, fast, and dependable file system for your users.


Disk Administrator

Disk Administrator (shown in Figure 8.1) is your tool for managing the hard disk subsystem, which includes any removable hard drives, such as Bernoulli, ZIP, and Syquest drives. You can use it to create, modify, or remove partitions; to format drives; to assign hard drive letters to specific volumes; and to create fault tolerant arrays of hard drives that will protect you from drive failure. If your Disk Administrator screen doesn't look like the one in Figure 8.1, it's probably just showing a different view. Choose Disk Configuration from the View menu to switch to the Disk Configuration view.

**FIGURE 8.1**

Disk Administrator window showing disk configurations

Any time you're in Disk Administrator, you can point at a portion of the disk with the mouse cursor and right-click to open a menu. When you right-click in an existing partition, you'll see this menu:



When you right-click on an area of free space, you'll see this menu:



Both menus provide shortcuts to many of the functions described later in this chapter.



CHAPTER 11

Printers and Other Resources

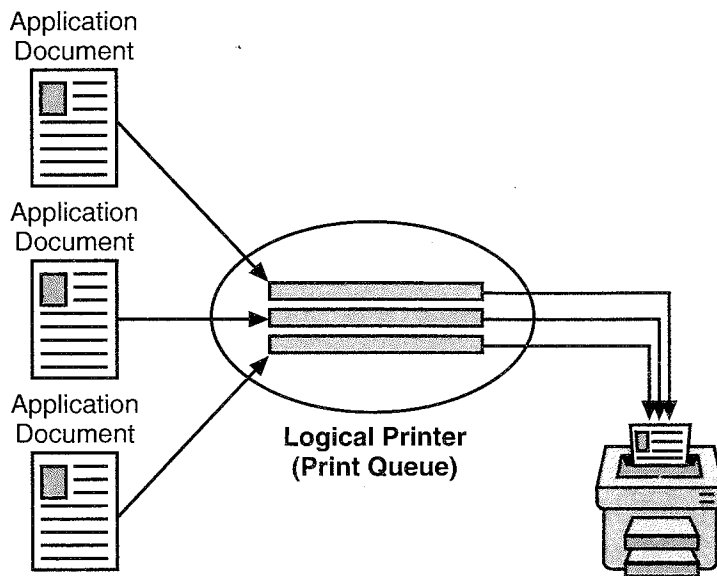
Hardware is expensive and companies don't like to invest in it until they absolutely have to. And as hard as it is to get permission to buy new equipment, it's even harder to get permission to buy equipment that is used only infrequently. From an economics point of view, the investment in equipment pays off only to the extent that the equipment is used. One of the great advantages of a network—although far from being the only one—is the ability to share equipment that otherwise would stand idle much of the time. Printers are a perfect example of this, and they are an obvious item to share. One could, certainly, consider providing an inexpensive dot matrix or ink jet printer for most users, but it's unlikely that many people would consider putting an expensive, high-end color printer or even a good, fast laser printer on everyone's desk. And you don't really need to. By putting one high-end printer on the network, you spread the cost and machine use across all of your users, which makes the investment in the equipment much easier to justify.

NOTE

Microsoft Windows NT version 4 uses some special and sometimes confusing terms when referring to printers. First it's important to differentiate between a print device, which is the actual machine that does the printing, and a printer, which in Microsoft terminology is the software interface between the application and the actual print device. Therefore, a printer in Microsoft terminology is actually a logical entity.

Printer Setup Options

To keep confusion to a minimum, we will refer to the print device simply as the "printer." We will refer to the software interface as the "logical printer." You should be aware, however, that the Windows NT documentation either doesn't use these terms or doesn't use them in this way. In a Novell NetWare or OS/2 networking environment, the term *print queue* is used instead of printer (meaning the logical printer), but the net effect is the same. Windows NT Server supports a broad range of printers. Figure 11.1 shows the simplest possible arrangement—a print job in Windows NT going to a logical printer, from which the job is spooled to a printer.

**FIGURE 11.1**

Example of jobs routing through a logical printer to a printer

You can have one logical printer associated with a single printer, which is the arrangement shown in Figure 11.1. Or you can have several logical printers associated with a single printer. In this arrangement, logical printers can be configured at different priority levels, so that one is for normal printing and the others are for jobs that can wait to be printed later. For a printer that uses both Postscript and PCL, having two logical printers allows users to choose either type of printing.

You also can have a single logical printer associated with multiple printers. If all of the printers use the same printer driver—an arrangement called a *printer pool*—a single logical printer will send jobs to the first available printer. The advantage of a printer pool is that the administrator can add or remove printers without affecting user configurations because the printers are interchangeable. The disadvantage of a printer pool is that there's no way to predict which printer will receive which job. So don't pool printers when they are physically far apart!

Planning Network Printing

As for everything you do on a network, you need to actually *plan* where and how your printers will be set up, configured, shared, and managed—a nuisance, but a necessary nuisance if you want to keep your trouble and your support calls down. You need to think about how your users really use printers, where the heaviest users are physically located, where to physically locate highly specialized printers such as plotters and color laser printers, and how you're going to physically connect all of the printers to the network.

Let's look at the last of these questions for a moment. To the average PC user, printers are always attached to a parallel port. But that doesn't work well for a network or for a server. Parallel ports require a fair amount of CPU attention to do their thing, which is the last thing you want on a server. You're also usually limited to three parallel ports. The other two choices are a serial connection (preferably using a "smart," multiport serial card) or a network connection. Either works well, but, in most cases these days, your best bet is a direct network connection.

How you choose to connect your printers will be influenced by several factors. You'll need to consider your physical layout. Are your users all in central physical locations? If so, you might find it easiest to simply recycle an older PC as a print server by installing a multiport serial card in it and using it to drive the printers. Or are your users (and their associated print needs) located in separate offices spread over several floors? If so, you'll probably want to use network connections, either external connections or connections built into the printers, to connect printers that are conveniently located for each group of offices.

There are two basic methods for connecting your printers directly to the network. You can use a high-end printer that comes with a network card that is either built in or available as an option. Or you can use a stand-alone network print server—the Hewlett-Packard JetDirect EX is a good example—that supports a variety of protocols and usually comes with drivers to support many network operating systems, including Windows NT Server. This is useful in a typical organizational environment, where your network might well consist of multiple operating systems—all of which require access to that expensive color laser printer.

In any case, once you've decided where and how to physically locate and connect the printers, you'll need to create and manage the logical printers that your users will actually see.

Managing Printers

Managing printers is not about the printers themselves but about how they are connected to and managed as part of your overall Windows NT Server network. You'll have to choose which networking protocol to use if the printers are network printers, decide which server you will use to manage them, and decide what functions you will allow your users to have access to.

Installing the DLC Protocol

If you have a Hewlett-Packard (HP) network printer, you have several options for controlling it; the simplest option is probably to use the DLC protocol. Another option is to use TCP/IP, which is probably the better option in the long run; but we'll



CHAPTER 14

Applications Software

One of the most confusing areas when dealing with the design of an enterprise network is the use of software applications over the network. Often even more confusing is the issue of software licenses. In this chapter, we will provide you with some guidelines to ease you through these murky waters and to help you prevent the errors often associated with using software applications over the network.

Windows NT Server as an Application Server

Microsoft Windows NT Server version 4 is a flexible network operating system that can be used in different ways:

- ◆ As a simple data file server, with clients storing their data in a central location.
- ◆ As the underlying operating system for a database server, such as Microsoft SQL Server 6.5, which runs as a service on Windows NT Server, storing the database and processing queries from clients.
- ◆ As an application server, providing the ability to run applications from the server either in a simple file serving mode, as older networked operating systems have run them, or in an actual client/server mode, where processing is shared between the client and the server, each taking responsibility for a portion of the application.

A key difference between Windows NT Server version 4 and other network operating systems is the ability of Windows NT Server to provide a stable, dependable server platform for client/server applications. The use of Windows NT Server as a full client/server applications server is what permits other tools, such as the remote administration tools, to do their jobs. In the case of remote administration tools, the applications and services that actually are being administered reside on the server; but the local applications that communicate with and control these services are run from a client machine, just as if they were being run locally off the console. With the addition of the Distributed Component Object Model (DCOM)—sometimes referred to as Network OLE—to Windows NT Server version 4, we expect to see many more applications take advantage of the Windows NT ability to run distributed applications.

Running Applications from the Server

The most common use and configuration of a Windows NT Server is that of a simple file server architecture, where the server is used to run applications and provide basic networking services. (See Figure 14.1.) Shared applications are stored on the central file server, and any client on the network can have access to them, even remote clients—via RAS.

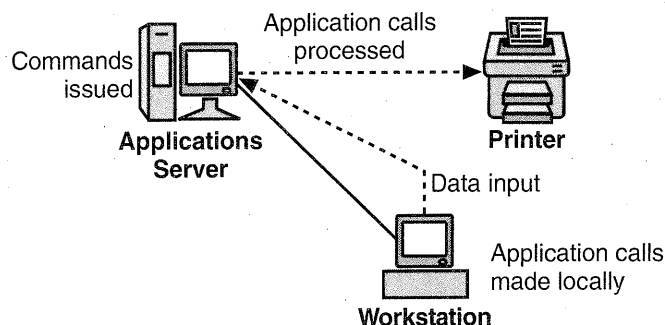


FIGURE 14.1

Process diagram for applications stored and run from the file server

The applications on the file server can be one of two types: a network version of the application or a stand-alone version of the application. Network versions of most applications, while often still functioning as straightforward file-shared applications, are moving more and more in the direction of running in full client/server mode, with the majority of the processing occurring on the server. Existing stand-alone applications might be stored on the server with client access through the server, but the actual processing is done on the client itself. (See Figure 14.2.)




FIGURE 14.2

Process diagram for applications stored on the server and run on the workstation

Both network and stand-alone applications have their advantages and disadvantages. Network versions are easy to monitor for access by the licensed users of the application. Applications run from a server might run more slowly because they are sent over the network, but this overhead can be more than offset by the potentially increased power of the server hardware and the more powerful database or application engine running there. It also might mean that you wouldn't need as powerful a workstation computer. With a single-user version of an application, you are likely to need more power at the workstation because all of the processing will be done there. The biggest problem for stand-alone versions of applications is in the area of license issues. If you own only 10 licenses of an application, that means that you can allow only ten users legal access to that application at any given time. Some applications leave this to the honor system by not restricting access to anyone. Other applications, however, do restrict access by checking users that log on against a license file associated with the application.

Either method for running applications from the file server will require that hard drive space be used for cache files and that RAM be used in the processing of the program. Running the application on the workstation rather than on the server will reduce the burden on the RAM requirements of the server, but it usually increases disk traffic.

Running Applications as a Service

The server components and routines that are a key part of the network operating system can use a great portion of the network server's memory and system resources. As Windows applications in general, and network server applications specifically, grow in size and resource requirements, it is essential that basic networking functions be provided in a way that utilizes memory and space efficiently.

In Windows NT Server version 4, most of the underlying network applications and functions run as a service. (See Figure 14.3.) As a service, the application runs in the background and in a minimized mode that makes calls for more RAM and resources only when it is called upon to perform its functions.

Windows NT Server as a Replication Server

It can be a complicated matter to keep versions of files straight on a network. Windows NT Server includes a replication service that ensures everyone is working with the same information. Directory replication also can help balance a server workload. When you have too many people who need access to a particular directory, you can export the directory to another server and point some of the workstations there. Directories are *exported* by the replication service on the server and dynamically updated when the master copy on the exporting machine is changed.



CHAPTER 19

Internet Information Server

Internet technology is not only the fastest growing area of computer technology, it's also the area of technology that is changing most rapidly. The same technology that is transforming the Internet is also providing the basis for "private networks" known as intranets. Armed only with the features included with Microsoft Windows NT Server version 4, you can create connections to the Internet, develop your own intranet, or both. In this chapter, we'll cover the Windows NT Server internetworking features and some of the ways you can use them.

Windows NT Server includes the Internet Information Server (IIS), a full-feature Internet server. With IIS, you can construct a World Wide Web server to display documents, graphics, sound, and even full-motion video to the world. In addition to providing these simple Web services, IIS provides ftp hosting services and a Gopher server, and it can link to a back-end database server to provide dynamic information as required. IIS also includes support for the Secure Sockets Layer (SSL) when a secure, encrypted connection is required.

The demand-based publishing technology that makes the Internet and World Wide Web such powerful and easy-to-use tools provides the same features and interface for internal corporate data in an *intranet* as opposed to an *internet*. Here, the tight integration of IIS with Windows NT Server provides an even more powerful tool than a traditional network, with easy links to Microsoft SQL Server and full Microsoft Windows NT security.

See Chapter 4, "Security Planning," for additional information on Windows NT security.

When you use IIS as both an Internet and intranet server, you can perform a variety of tasks. For example:

- ◆ Publish a Web *home page* that provides current information about your department to the rest of the company, or information about your company to the rest of the world.

- ◆ Run a business, take orders, and create an online catalog.
- ◆ Create client/server applications that provide real-time information about current production numbers to other departments.
- ◆ Create a repository of up-to-date files that enable customers to download patches, updates, or even complete software programs using the ftp server.
- ◆ Publish older, plain-text documents from a variety of sources using the Gopher server.

IIS Installation

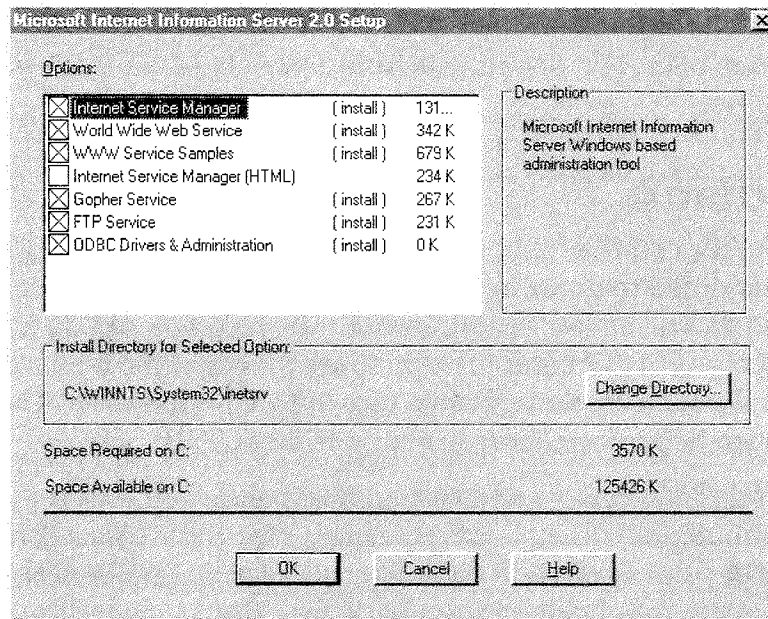
You were offered the choice of installing IIS when you initially installed Windows NT Server. If you opted to do so, by all means skip ahead to the section “IIS as an Intranet Server.” However, if you didn’t install IIS then, you can do it easily now. You’ll need to have TCP/IP configured before you begin the IIS installation, of course, but you should have that all taken care of by now. If not, see Chapter 7 and Chapter 17 for detailed information on configuring TCP/IP.

There is one important security matter to consider before you start the installation. You probably should install the home directories for the various services on an NTFS volume. The security of a FAT volume just isn’t sufficient for running this sort of access, even on an entirely internal network. In addition, if you’ll be using IIS as an external Internet server that is connected directly to the Internet, we strongly advise you to provide an additional layer (or layers) of security or a firewall between the server and your internal network.

When you installed Windows NT Server initially, an icon was added to the Administrator’s desktop for running the IIS installation program. Double-click this icon to begin the installation. (You also can run the installation program from the Add/Remove Programs application in Control Panel.) Whichever way you decide to run the installation, you’ll be presented (as shown in Figure 19.1 on the following page) with a list of components to install:

- ◆ **Internet Service Manager** The administrative tool for managing your IIS server
- ◆ **World Wide Web Service** The main Web server program
- ◆ **WWW Service Samples** Sample HTML, database, and ActiveX files
- ◆ **Internet Service Manager (HTML)** An HTML version of the administrative tools
- ◆ **Gopher Service** A publishing server that supports the hierarchical, distributed, text-based, Gopher service

- ◆ **ftp Service** A publishing server that supports the distribution of files using ftp
- ◆ **ODBC Drivers and Administration** Open Database Connectivity drivers that support connection to a database

**FIGURE 19.1**

The Internet Information Server installation screen

NOTE

If you have already installed Microsoft SQL Server 6.5, do not use the IIS ODBC SQL Server drivers. The drivers that ship with SQL Server 6.5 are more recent. If you do install the ODBC SQL Server drivers that ship with IIS, you can reinstall the ODBC SQL Server drivers that ship with Microsoft SQL Server 6.5.

By default, the HTML version of the administrative tool set is not installed. If you will be administering your server remotely, you might find the HTML version of Internet Service Manager a useful tool. Otherwise, you'll probably find it preferable to use the native Windows tools.

The default directory for IIS is %SYSTEMROOT%\SYSTEM32\INETSRV, although the default for the actual publishing services is under \INETPUB on the system drive. If you want to run your IIS installation in its own disk volume, now's the time

to install it on its own disk volume. We tend to favor leaving the IIS program files in the default directory and then installing the actual publishing services on their own disk volume, where this is possible. Separating the IIS program files and the publishing services makes it easy to segregate the publishing services from the rest of your server's functions. However, if the sole function of the server is to act as your Web server, you might as well stick to the default directories.

Select the options you need, and click OK. The dialog box shown in Figure 19.2 will appear. From here, you can select the root directories for the various publishing services. Spend a few moments now thinking about where you want to install the root directories because it can be a nuisance to move the publishing services later.


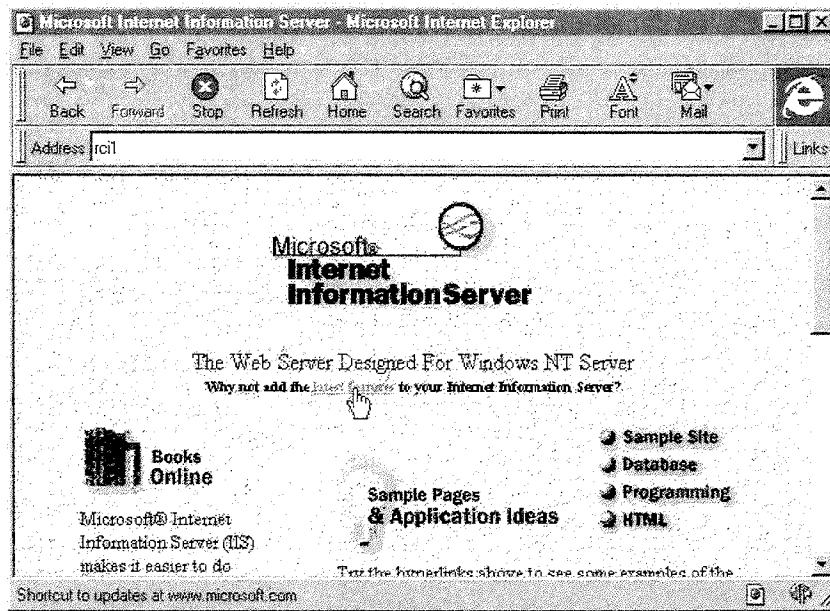


FIGURE 19.2

IIS publishing service root directories

After you have selected the location for your root directories, click OK to do the actual installation. The program will ask you whether you want to create new directories if they don't already exist. Click Yes to create them. The installation program will install the services you have selected and start them up. If you select ODBC Drivers and Administration, the program will prompt you for the available ODBC drivers. By default, the only ODBC driver is the one for SQL Server. If you've already installed SQL Server 6.5, do *not* install the IIS ODBC SQL Server driver; use the ODBC driver that ships with SQL Server 6.5.

The final step in the installation process is to test the server. From another computer on your network, start Internet Explorer and connect to the server you just set up. Because this is an internal network connection, you can shortcut the process by typing the name of the computer on which you just installed IIS in the address text box as shown in Figure 19.3 on the following page.

**FIGURE 19.3**

Using the short version of the URL on an intranet

By default, IIS supports NetBIOS names as a shortcut for the home page of servers on an intranet. The full version of the URL (pronounced "earl"), however, would be

http://<fully qualified domain name>

If you have a problem connecting to the server, try using the server's IP address instead of its name. This will narrow down the source of the problem; if you can connect with the server IP address but not with the server name, you will have to troubleshoot your name resolution subsystem. Start at the WINS server. If that appears to be OK, check the DNS server.

IIS as an Intranet Server

The term *intranet* is relatively new and means the interconnected computers within an organization, regardless of operating systems, physical connection type, and so forth. This may be a simple Windows NT network with a half dozen workstations all connected together and all running the same operating system or a complex, world-wide enterprise network that runs a wide variety of operating systems connected via a variety of permanent, semipermanent, and intermittent networks using a variety of network protocols.

What distinguishes an intranet from the Internet is that the computers on an intranet are all part of a single organization. Furthermore, we generally assume that they are logically or physically isolated from the Internet. This isolation can be purely physical, with no portion of the network actually connected to the Internet—clearly the safest way to isolate your internal network, or the isolation can be a “logical” isolation—a firewall of some sort that permits legitimate use of the Internet from within the organization but shields the internal network from unauthorized external users.

As more and more organizations and corporate users experience the power and ease of use of the Internet, there is more pressure to leverage these same technologies to create an intranet within the organization. An intranet can provide an effective and efficient communication tool within an organization. Because it is a purely demand-based method of communication, it keeps network traffic to a minimum. Only those users who actually need a particular piece of information will download it and, then, only when they need it. The alternative—a traditional push system of updating files, memos, and information—requires that everyone who might possibly need a particular piece of information must get it sent to them, which is a much less efficient use of network resources.

IIS is ideally suited as an intranet server for two reasons:

- ◆ It supports a variety of protocols and programming languages.
- ◆ It takes advantage of the native Windows NT security features, which allow you to control access to sensitive areas of information easily at the same time you make the information readily available to those who need it.

IIS supports the ActiveX standard created by Microsoft. (Microsoft intends to transfer the standard to an independent group.) ActiveX allows software components to interact with client computers from a variety of operating systems without regard to the language or the operating system with which the object was created. ActiveX components can be:

- ◆ **Controls** Software components that run Web pages and provide interactive and user-controlled functions. ActiveX controls allow users to view animation, audio, and video without opening separate programs.
- ◆ **Scripting** A standard that provides a generic way for clients to execute scripts written in any scripting language, such as Visual Basic Script and Javascript.
- ◆ **Documents** An extension of OLE compound documents that supports using Web documents to open programs. It allows you to use non-HTML documents such as Excel or Word documents from within a Web browser such as Internet Explorer.

IIS also supports the Internet Server Application Programming Interface (ISAPI), which provide for significantly faster and more versatile connections than the Common Gateway Interface (CGI).

ftp Service

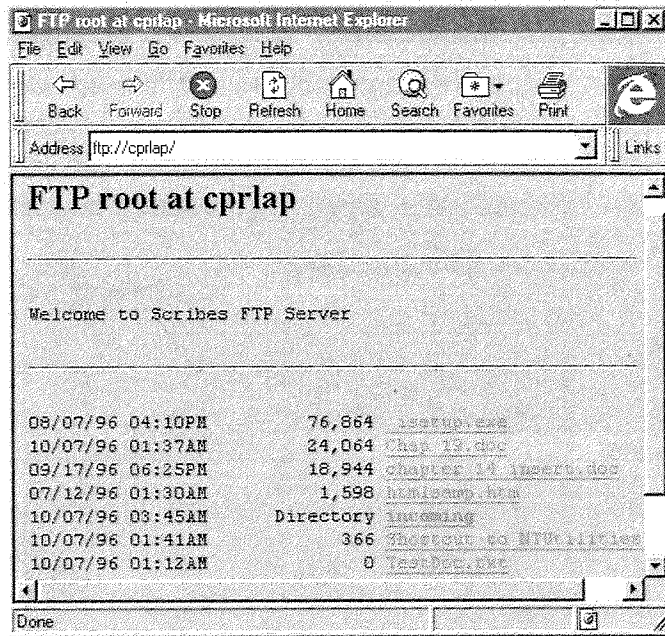
Ftp service provides a platform-independent way of transferring files between computers. There are ftp servers and ftp clients available for virtually every operating system and hardware platform made in the last 20 years. IIS supports both anonymous ftp and traditional ftp, although it has a strong preference for anonymous ftp—and with cause. One of the problems with traditional, account-based (or user-based) ftp is that it requires you to send your password *unencrypted*. This is a “bad thing,” obviously, although it might be perfectly acceptable in an intranet situation.

Anonymous ftp, on the other hand, assumes that everyone should have access to the files and that no user password information is actually passed. Modern Web browsers, such as Internet Explorer and Netscape Navigator, allow you to connect directly to an anonymous ftp server just as if it were another Web page. With IIS, all you have to do to make files available for downloading is place them in the FTPROOT directory. Figure 19.4 shows a simple example of a connection to an intranet ftp site using Internet Explorer and IIS. Type the URL in the text box in the form of `ftp://<servername>`, and you’re in business. Any files in the directory become visible and available for downloading simply by clicking them. Meanwhile, Windows NT Server protects the rest of your server’s file system by hiding any directories above the FTPROOT directory, masking them from prying eyes.

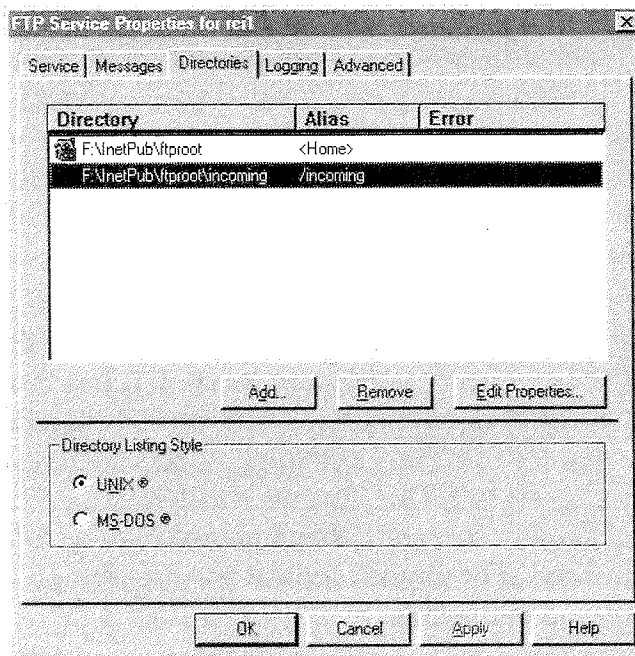
IIS also can be configured to permit uploads. The default is to have this set to “off,” but you can change this by running Internet Service Manager and double-clicking the ftp server you want to change. Click the Directories tab to bring up the dialog box shown in Figure 19.5.

From here, it’s easy to set a directory to allow uploads as well as downloads. It’s a good idea to create a special directory for uploads separate from the regular download directory. Files are allowed to be uploaded only to this directory and are checked carefully for viruses and other problems before being made available to the rest of the network.

To make an upload directory, click Add and type the name of the directory in the Directory text box. If you haven’t already created the directory, click Browse to locate where you want to put it and type the name of the directory in the New Directory Name text box. After you’ve added the directory, all you have to do is give the directory Write permission. For added security, you should remove the Read permission from the directory; this makes it impossible for users to see what’s in the directory but they can still upload files to the directory. This configuration is

**FIGURE 19.4**

Example of a connection to an ftp server

**FIGURE 19.5**

Using Internet Service Manager to set ftp directories

shown in Figure 19.6, where we've created a special "incoming" directory for users uploading files to our server. The files in this directory cannot be seen because the Read permission isn't set, but users can upload files to it. The administrator then can check any uploads carefully before moving them into public view.

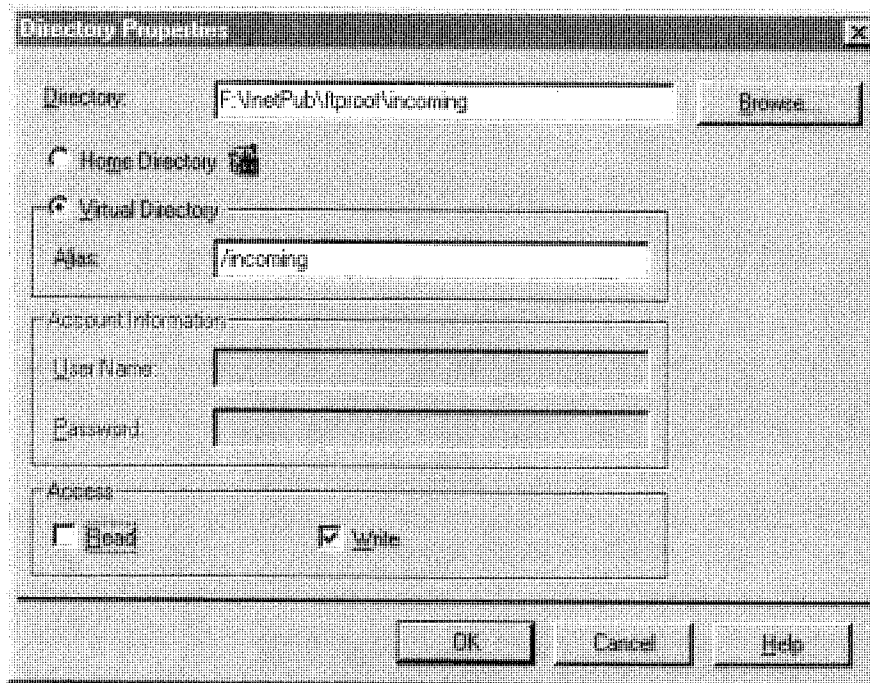


FIGURE 19.6

Example of an upload-only directory for the ftp server

Gopher Service

Gopher service is an older, text-only service that allows you to create a hierarchical, distributed database of text files that can be searched, viewed, and downloaded, without having to worry about where the files are located. It has been, in many ways, supplanted by the World Wide Web service; but it still has merit when you are searching, viewing, or downloading large amounts of text documentation that would be cumbersome to convert to HTML format. The federal government, for example, runs a number of Gopher servers. Other sources of useful Gopher documents are academic institutions. For a quick look at the breadth of information available for Gopher clients, open the URL:

<http://galaxy.tradewave.com/GJ>

This is a repository of some of the jewels of the "gopherspace," as the collection of interconnected Gopher servers is known.

Microsoft Windows NT Server 4.0

APPENDIX G

If You Design Networks with Windows NT Server 4.0, Think of This as Your Most Important Component.

In this practical, one-volume handbook and reference, system administrators and managers will find powerful help with the networking challenges they face most often. It's all here—planning the right network design; installing, tuning, and maintaining the system; and recovering from disaster.

And unlike books that are little more than presentations of system capabilities, *RUNNING MICROSOFT WINDOWS NT SERVER 4.0* is a comprehensive road map. The emphasis is always on planning, strategy, and the needs of your organization—so you always know where you're going and why.

If you work with Windows NT Server 4.0, this book is for you. It's complete enough for the technically advanced, yet it's friendly and accessible. A quick introduction to basic concepts helps you if you're developing a Windows NT network from scratch. And hundreds of pages of practical, hands-on guidance from these experienced professionals help you implement a system that's tailored to your organization. You'll also learn how to work with Microsoft Exchange and with Internet server tools such as Microsoft Internet Information Server.

Get *RUNNING MICROSOFT WINDOWS NT SERVER 4.0*. It fits perfectly between Microsoft Windows NT Server Resource Kit and Microsoft's official multivolume training materials. And that makes *RUNNING MICROSOFT WINDOWS NT SERVER 4.0* the perfect handbook to use every day.

Windows NT/Networking




7 90145 13331 1

ISBN 1-57231-333-1



9 781572 313330

9 0000 >



U.S.A. \$39.95
U.K. £36.99
Canada \$54.95
[Recommended]

www.microsoft.com/mspress/



Start Faster and Go Farther with Help from Microsoft Press.

Whether you're a beginner, a veteran, or a power user, Microsoft Press has books to fit your needs and your style.



Select Editions—

Comprehensive information in one volume




At a Glance series—

Easy, visual information, just when you need it



Step by Step series—

Microsoft's self-paced training kits



Field Guides—

Compact quick references



Starts Here™ CD-ROM series—

Interactive training

Microsoft® Press

Makes it easy to set up and use a small network in your home or office!


APPENDIX H



The Little

Network Book

for Windows and Macintosh



*A friendly introduction
to the basics of running
a small network, including
setup, Internet access,
sharing printers and files,
troubleshooting, and more!*

LON POOLE & JOHN RIZZO

ILLUSTRATED BY JOHN GRIMES

ROKU EXH. 1002

The Little Network Book

Lon Poole and John Rizzo

Illustrations by John Grimes



Peachpit Press
Berkeley ▼ California

APPENDIX H

Peachpit Press

1249 Eighth Street
Berkeley, CA 94710
510 524 2178
fax 510 524 2221
Find us on the Web at <http://www.peachpit.com>

The Little Network Book

Copyright © 1999 by Lon Poole and John Rizzo
Cartoon illustrations copyright © 1999 by John Grimes
grimescartoons.com

Peachpit Press is a division of Addison Wesley Longman

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the publisher. For information on getting permission to reprints and excerpts, contact Gary-Paul Prince at Peachpit Press.

Notice of Liability

The information in this book is distributed on an "as is" basis, without warranty. Although every precaution has been taken in the preparation of this book, neither the author nor Peachpit Press shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the instructions contained in this book or by the computer software and hardware products described herein.

Trademarks

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Peachpit Press was aware of a trademark claim, the designations appear as requested by the owner of the trademark. All other product names and services identified throughout this book are used in editorial fashion only and for the benefit of such companies. No such use, or the use of any trade name, is intended to convey endorsement or other affiliation with the book.

Editors: Lisa Theobald, Judy Ziajka
Production Coordinator: Amy Changar
Compositor: Owen Wolfson
Interior Design: Robin Williams
Cover Design: **TMA** Ted Mader Associates
Cartoon Illustrations: John Grimes
Indexer: Karin Arrigoni

ISBN 0-201-35378-4


0 9 8 7 6 5 4 3 2 1

Printed and bound in the United States of America

Share Printers

The desire to share a high-end printer among several computers is reason enough to install a network. Printing over a network is a lot like printing to a printer connected directly to your computer. But before you can use a shared printer, you'll need to install and configure some software.

As with other aspects of networking, sharing a printer among Macintosh computers is quite a bit different from sharing among Windows machines. In some cases, however, you can set up a network so that both PCs and Macs can use the same printer.



Types of Printers to Share



Local Printers

On a network, you can share just about any type of printer, from an inexpensive inkjet to a big, fancy, feature-filled laser printer. Some printers are designed to connect directly to an Ethernet or LocalTalk network; these **network printers** are made specifically for sharing. But you can also share many of the **local printers**, also known as **personal printers**, that directly connect to one of the computers on the network. Where a printer is connected to the network affects how you set it up for sharing but not how you print to it. We'll detail all these procedures in this chapter.

The Local Angle: A **local network** is the network in your home or office, but a **local printer** is plugged directly into one computer. You can set up your network to share both local and network printers.

Local printers are designed to **connect to one computer**. Most inkjet printers are local printers, as are the older dot-matrix printers that some people still use for printing on carbon-copy paper. Some laser printers are local printers as well.




When you share the local printer connected to your computer, your computer provides the network connection to other computers. Other network computers that want to print to your shared local printer must contact your computer over the network and send your computer the pages to be printed. Your computer has to do a bit of processing before it tells its printer to print the pages. This means your computer must be up and running if anyone wants to use your shared local printer, and it also means that processing some of those printing jobs could cause some minor delays on your machine if you're simultaneously working on your computer as another user is printing.

An independent network printer isn't "local" because it doesn't connect to any single computer. Rather, an independent printer **connects directly to the network**, just like a computer (via a hub, for example). In fact, an independent network printer actually is a computer—with a processor and RAM of its own. This additional equipment enables the network printer to communicate directly with the computers on the network.

An independent network printer is always available to all network computers, just as if it were connected directly to each computer. The biggest drawback to this setup, however, is that independent network printers are usually more expensive than local printers.

Independent Network Printers

Many laser printers are independent network printers; most inkjet printers are not.



Ports and Adapters

Independent network printers are easy to identify. Just look for the network port—either a **10BaseT Ethernet** port or a **LocalTalk** port. (Review Chapter 2 for details on Ethernet and LocalTalk.)

*Network adapters made for printers are also called **print servers**.*

If a printer doesn't have a network port, you may be able to add one by installing a **network adapter** card in the printer. Adding an internal or external network adapter essentially turns a local printer into an independent network printer.

To determine whether your printer has an expansion slot for a network adapter, check the printer manual or ask the printer manufacturer. Even if your printer doesn't have an expansion slot for an internal network adapter, you may be able to use an external network adapter to connect the printer to your network. The printer manufacturer may also be able to tell you whether there is an external network adapter that will work with your printer, or you can check with a computer store or catalog.

Mind Your Network Protocol

Some older network printers have only a LocalTalk port—often called an **AppleTalk port** by printer manufacturers, because LocalTalk can use only the AppleTalk protocol. Printers with an Ethernet port can use many network protocols, including AppleTalk and TCP/IP.

AppleTalk is chiefly a Macintosh protocol. PCs can't communicate with network printers that use only the AppleTalk protocol because Windows 95 and 98 don't normally support it. For example, a printer with a LocalTalk port but no Ethernet port can use only AppleTalk.

You can, however, add AppleTalk support to Windows 95 and 98 by installing additional software that's not included with Windows. Alternatively, you may be able to bypass the AppleTalk protocol altogether by connecting the printer to a PC as a local printer. These solutions are all more thoroughly described in "Connecting PCs to AppleTalk Printers" later in this chapter.

Getting Ready to Share a Printer

To share a printer among the computers on your network, that printer needs a **physical connection**—either to one of the network computers if it's a local printer or directly to the network if it's a network printer. Each network computer must have installed the appropriate **driver software** for the printer as well as **network printing software**.

A local printer connects to one of the ports on the back or side of one of your computers. Generally, you connect the printer to the computer using a cable included with the printer. The type of cable depends on whether the printer was made for the PC, Macintosh, or both.

- Most local printers for PCs connect using a **parallel cable**.
- Most local printers for Macs connect with a **serial cable**.
- Newer printers can connect to both PCs and Macs using a **USB cable** (universal serial bus).

We'll explain how to connect a local printer later, in "Connect the Printer to the PC" and "Connect the Printer to the Mac."

USB Converters: If you want to connect a local printer that doesn't have a USB port to a computer that has only USB ports, you may be able to buy a converter cable for the printer. Hewlett-Packard and Epson sell USB converters for their own printers, and other companies also offer converter cables. Check with your local computer store to see what's available.

Like a network computer, an independent network printer connects directly to your network. If you have an Ethernet network, you'll need 10BaseT patch cable long enough to reach a nearby network jack or your Ethernet hub. If you have a LocalTalk network, you'll need a LocalTalk connector and a phone cord. All this equipment is described in Chapter 2.

If you've set up an Ethernet network and your network printer has only a LocalTalk port, don't panic. You can get a LocalTalk-to-Ethernet converter. These are described in the later section "Connect a LocalTalk Printer to an Ethernet Network."

Whether you have an independent network printer that connects directly to your network or you will share a local printer that connects to one of your computers, all of your computers must have compatible **driver software** for that printer to use it via the network. The driver software communicates with the printer using a set of commands, a kind of **printer language**. Using the printer language, the driver tells the printer what to print on each page, along with instructions on how many copies to make, what paper tray to use, and which special features of the printer to use.

Local Printer Connection



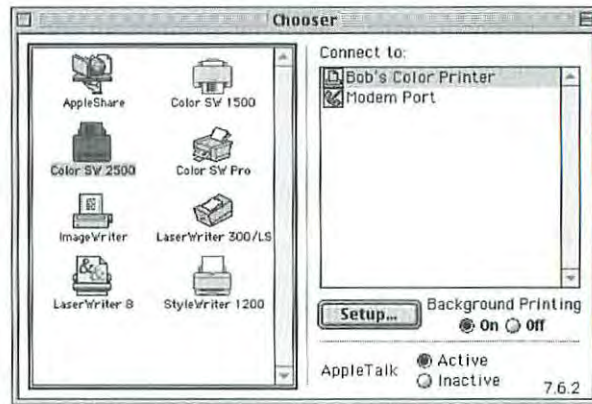
Network Printer Connection

A network printer may have its own network port, or it may plug into an external network adapter.



Drivers and Printer Languages

After you enable sharing of a local printer, the Chooser identifies it by the name you gave it rather than by the port to which it's connected.



The Chooser lists a shared printer's name (not its port).

Setting Up Independent Network Printers

Independent network printers are easier to set up than shared local printers, because independent printers are inherently shared devices. You simply plug the printer into your network, perhaps assign it an IP address, and then it's available to all the computers on your network. You don't have to configure a computer to act as a host for an independent network printer.

Connect the Printer to the Network

Network printers have an Ethernet port, a serial port for a LocalTalk connector, or both. These ports are just like the ones you'll find on your computers, so you can plug a printer into your network by following the procedures described in Chapter 2 for connecting computers to the network. Here's a recap:

- **For Ethernet:** Run an Ethernet patch cable from the printer's Ethernet port to your Ethernet hub or to a nearby Ethernet wall jack.
- **For LocalTalk:** Plug a LocalTalk connector into the printer's serial port, and use a phone cord to attach the printer to your LocalTalk daisy chain network.

Using an Internal or External Adapter

If the printer doesn't have a built-in network port, you'll have to install and configure the internal or external network adapter, or print server, that was described earlier in this chapter. Installing an internal print server usually means inserting an adapter card into the printer's expansion slot. Installing an external print server usually involves connecting

a cable from the printer to the print server and plugging the print server into a power outlet. After installing the internal or external print server hardware, you may have to install some software on each computer that's going to use the printer. Follow the detailed instructions included with the print server.

In some printers, including some HP LaserJet models, you can replace a LocalTalk adapter card with an Ethernet adapter card. However, you can't add Ethernet to Apple LaserWriter models that don't already have it.

Still, there are other ways to connect a LocalTalk printer to Ethernet. One method is to use an external, self-contained LocalTalk-to-Ethernet converter box such as Farallon's iPrint LT. You connect the converter to your printer's LocalTalk connector with a phone cord, and you connect the converter to your Ethernet hub or to a nearby 10BaseT wall jack with an Ethernet patch cable. The iPrint LT lets you connect not just one LocalTalk printer but a daisy-chain of up to eight LocalTalk devices to one port on your Ethernet hub.

Another solution is Apple's free utility software, LaserWriter Bridge, which you install on a Macintosh that has both LocalTalk and Ethernet connections. LaserWriter Bridge is included with some of Apple's LocalTalk printers, and it is available from Apple's Software Updates library on the Web at <http://www.info.apple.com/swupdates>.

A Macintosh running LaserWriter Bridge must be turned on for computers connected to the Ethernet network to access the LocalTalk printer.

Here's how to set up LaserWriter Bridge:

1. Using LocalTalk connectors and cables (as described in Chapter 2), connect the LocalTalk printer to a Macintosh on the Ethernet network.
2. Install LaserWriter Bridge software on this Mac by dragging the software's icon from the installation disk to the Mac's System Folder icon (*not* to the System Folder window).
3. A message asks you to approve copying LaserWriter Bridge to the Control Panels folder. Click OK.
4. Open LaserWriter Bridge from the Control Panels folder on your hard disk. (*Don't* open the LaserWriter Bridge on the installation disk.)
5. Click the On button in LaserWriter Bridge.
6. Restart the Macintosh to make the connected LocalTalk printer available to other computers on the Ethernet network.

Connect a LocalTalk Printer to an Ethernet Network

Using LocalTalk-to-Ethernet Hardware

Using LaserWriter Bridge Software



Set Up LaserWriter Bridge

Set a Network Printer's Address



An independent network printer that uses the TCP/IP protocol must have its own IP address. The method of assigning an IP address to a printer varies from printer to printer and even among printers of the same make. You'll need to follow the instructions that came with your printer. Here we'll give you a general idea of how some printers handle IP addressing.

The rules for assigning IP addresses to computers are discussed in Chapters 3 and 4. The same rules apply to printers.

Usually, you set the printer's IP address using software running on a PC or Macintosh. Some printers require that you use a special utility program included with the printer. With other printers, you use a general-purpose Internet program, such as a Web browser.

- In some cases, you must first plug the printer into the computer's parallel port or serial port. Then you run the printer utility on the computer to set the printer's IP address. Finally, you unplug the printer from the computer and plug it into your Ethernet network.
- In other cases, the printer has a preset IP address, such as 0.0.0.0, so you can connect the printer to your Ethernet network right away. You have to change this IP address to be consistent with the other IP addresses on your network, but you can do this from any computer on the network, and you don't have to plug this printer into a computer to set the address.

Some printers can be set to get an IP address automatically from a server, such as a DHCP server, if your network has one. (Chapter 5 describes one way to add a DHCP server to your network.) With these printers, you can still manually assign a static IP address.

- Some printers that support TCP/IP will let you view the IP address on the printer's LCD display via the control buttons on the printer. Learning how to do this is a little like learning how to program your VCR; it's not always straightforward, so you'll need to follow the instructions in your printer's manual.

After setting up one or more shared printers on your network, you and other network users can print to them. All the shared printers that are available to Windows PCs on a network—including local printers shared by PCs and independent network printers that PCs can use—appear in the Network Neighborhood window of each PC. The Network Neighborhood shows you which printers are connected to your network, but your PC can't use a shared printer unless its icon appears in your PC's Printers folder.

If your Printers folder already contains an icon for a shared printer that you want to use, you're all set. You can begin printing to this printer whenever you like. If you need to add an icon for a shared local printer or an independent network printer, you can do so from the Network Neighborhood or from the Printers folder. If you're not sure whether you need to add an icon for a shared printer, go ahead and add it. Having more than one icon for a printer in the Printers folder is okay.

The quickest way to add a shared printer's icon to the Printers folder is from the Network Neighborhood. Follow these steps:

1. Open the Network Neighborhood (by double-clicking its icon on the desktop).
2. Look for the icon of the printer whose icon you want to add to the Printers folder. If the printer is attached to one of the computers on your network, you'll have to double-click that computer's icon in the Network Neighborhood to see its shared printer.
3. Right-click the icon of the printer you want to add, and choose *Install from the shortcut menu*. The Add Printer Wizard starts.
4. When the Add Printer Wizard asks whether anyone prints from MS-DOS programs on this PC, answer appropriately, and then click the Next button.

If you need help on printing in general, open the Windows Help system (click Start and then click Help), click the Contents tab, and look under Print in the How To section.

Using Shared Printers from a PC

If you're using a Macintosh, not a PC, skip ahead to "Using Shared Printers from a Macintosh."

Install from the Network Neighborhood





The Network Neighborhood shows network printers that you can add to the Printers folder.

5. The Wizard may now display lists of printer makes and models. Select the network printer's manufacturer in the left-hand list and the printer model in the right-hand list.

If the printer isn't listed, you'll need to provide a disk that has the printer's driver software:


- Click the Have Disk button, and in the dialog box that appears, specify the location of the driver for the network printer. Then click OK to return to the Add Printer Wizard.
- If Windows asks for the disk you specified, insert it so that the driver software can be copied to the PC's hard disk.

The printer make and model lists in step 5 don't appear if you specified a shared local printer that is **online** (meaning this PC can communicate with the printer's host PC). In this case, the Wizard will copy the driver files from the printer's host PC.

6. If the Printers folder already contains another printer's icon, and the Add Printer Wizard will ask whether you want the printer you're adding to be the default printer in Windows applications. Click Yes or No. (The default printer should be the printer you use most of the time. You can change the default printer at any time.)
7. If the printer has been assigned a password, you must enter it before you can use the printer.



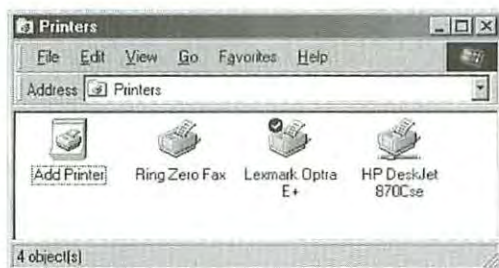
8. Name the printer when you're asked. This name will appear below the printer's icon in the Printers folder on your computer. The name will appear on your computer only, not on other computers.



Make a network printer the default printer if you will usually want Windows applications to print to this printer.

9. The Wizard asks if you want to print a test page. Click Yes, since this is usually a good idea, and then click Next.
10. Click Finish, and Windows copies files for the shared printer to the PC's hard disk.

A new icon for the shared printer now appears in the Printers folder. The icon has a network cable at the bottom to indicate that it's a shared printer. The default printer has a check mark on its icon.



An icon with a cable at the bottom is a network printer, and the icon with a check mark is the default printer.

Each network computer can have its own particular name for the same shared printer.

Change the Default Printer:

To make a different printer the default printer, right-click its icon and choose Set As Default from the shortcut menu. The "Default" check mark will now appear on the selected printer's icon.

Install from the Printers Folder



Instead of going through the Network Neighborhood to add a printer icon to the Printers folder, you can create the printer icon from **within the Printers folder** itself. This takes a little more work in the beginning because you have to specify the location of the shared printer, but the last part of both procedures is the same. Follow these steps:

1. Open the Printers folder in My Computer. Or, from the Start menu, select Settings and then Printers.
2. Double-click the Add Printer icon to launch the Add Printer Wizard, and then click the Next button to get started.
3. Select Network Printer and click the Next button.

Network Printer means a shared printer connected to another computer or an independent printer directly connected to your network.

4. The Wizard asks you to specify the network path. (A network path looks something like this: `\\computer name\share name`.)

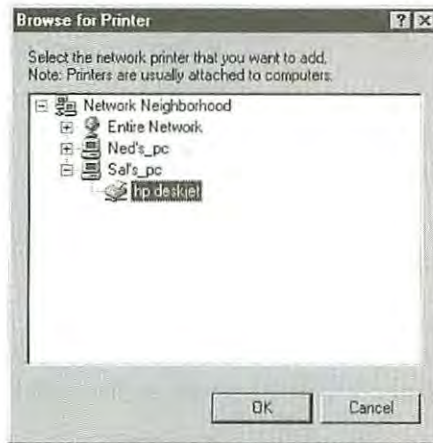


The network path tells Windows where to find the shared printer you want to use.

Browse for the
Printer's Location

5. If you don't know the printer's network path, click the Browse button. You'll see a list of the computers and independent printers on your network.
 - You can click a particular printer's icon to specify its network path. Or, to find a printer connected to another computer, double-click the computer icons in this list until you find the printer you want. Then select the printer icon.

- Click the OK button, and the network path for the printer you selected will appear in the Add Printer Wizard.



If you don't know the network path, you can look for it in this Browser dialog box.

6. The Wizard also asks you to indicate whether you or anyone on your network prints from MS-DOS applications. Answer Yes or No, and then click the Next button.
7. With the network path entered, click the OK button.
8. If the Wizard tells you the printer you specified is offline, check the network path displayed at the top of the Wizard to make sure you typed it correctly. If the path is wrong or you're not sure, click the Back button and repeat steps 4 through 6. If you're sure the path is correct, click the Next button to proceed.
9. From here, the procedure is the same as for creating a printer icon from the Network Neighborhood. Continue at step 5 of the earlier section, "Install from the Network Neighborhood."

*The network printer will be **offline** if the computer it's plugged into is turned off.*

Windows 95 and 98 don't include built-in support for the AppleTalk protocol, which some network printers use to communicate over Ethernet and LocalTalk networks. Many printers that use AppleTalk can also use other network protocols, including TCP/IP, but some printers can use only AppleTalk. These AppleTalk-only printers are designed primarily for Macintosh networks, because all Macs support AppleTalk. If you want to use an AppleTalk-only printer in a network that includes PCs as well as Macs, you can add AppleTalk capability to Windows 95 or 98 by purchasing and installing additional software.

Connecting PCs to AppleTalk Printers

Use PC MACLAN

One way to add the AppleTalk protocol to Windows is by installing Miramar Systems's **PC MACLAN** (<http://www.miramarsys.com>). Installing PC MACLAN on a PC makes AppleTalk printers appear in the PC's Network Neighborhood along with printers shared by other PCs. (PC MACLAN also lets Macs and PCs share files, as described in Chapter 7.)

As is true for any other printer, the PC with PC MACLAN must have appropriate driver software to use an AppleTalk printer. This is no problem for printers that use PostScript, and most AppleTalk printers do. Windows includes driver software for a number of PostScript printers, and additional driver software is available for Windows from printer manufacturers.

PC MACLAN does not enable printing to non-PostScript printers, such as a local StyleWriter that is shared by the Mac it's connected to. That's because Windows-based driver for these printers is not available.

As usual, you'll create an icon in the Printers folder.



To use an AppleTalk PostScript printer on a PC with PC MACLAN, you need only create an icon for it in the PC's Printers folder using the Add Printer Wizard (as described in "Install from the Network Neighborhood" and "Install from the Printers Folder").

If the Wizard asks you to select the printer make and model but doesn't list the correct combination, and you don't have a disk with Windows driver software for the AppleTalk PostScript printer, try selecting Apple in the list of manufacturers and LaserWriter II in the list of Printers. This driver may not provide access to all the features of the AppleTalk PostScript printer, but it may at least enable the PC to print on this printer.

Use DAVE

Another product, Thursby Systems's **DAVE** (<http://www.thursby.com>), takes a different approach to letting PCs use an AppleTalk PostScript printer. Installing this software on a Macintosh allows PCs on the network to access one AppleTalk printer that uses PostScript. The PCs see the AppleTalk network printer in the Network Neighborhood as a share of the Mac that's running DAVE.

To use an AppleTalk printer shared by a Mac with DAVE, you just add the shared AppleTalk printer's icon to the PC's Printers folder using the Add Printer Wizard, as described earlier in this chapter. If you can't find the Windows driver software for your particular AppleTalk PostScript printer, selecting the Apple LaserWriter II driver listed in the Add Printer Wizard may at least enable the PC to print.

Many AppleTalk printers have a parallel port as well as an Ethernet or LocalTalk port. You can plug one of these printers into a PC's parallel port, making it a local PC printer, and set up that PC to share the printer with other PCs. The local printer doesn't need to use AppleTalk or any other network protocol to communicate via its parallel port. You can keep the printer connected to your network so that Macs can still use it. The printer accepts print requests from PCs via its parallel port and from Macs via its network port.

If you have experience printing from a Mac, you'll find that accessing a shared printer via AppleTalk is similar to accessing a local printer. You indicate which printer you want to use, and it becomes the **default printer** (the active printer) to which all Mac applications print.

You can designate the default printer in the Chooser. Follow these steps:

1. Open the Chooser (from the Apple menu). On the left side of the Chooser, you'll see an icon for every printer driver installed on the Mac.
Some icons in the Chooser have nothing to do with printing. The **AppleShare** icon, for example, is used for connecting to shared files on other Macs (as described in Chapter 8).
2. Look at the bottom right corner of the Chooser and make sure AppleTalk is set to Active.
3. Choose the icon of a driver for the shared printer that you want to designate as the default printer. The Mac OS includes printer driver software for almost all printers made by Apple Computer, as itemized in the table earlier in the section, "Setting Up a Local Mac Printer for Sharing."

Some PostScript printers use the **AdobePS** driver or an older **PSPrinter** driver, or a PostScript driver based on one of these. Many if not all of these printers will also work with Apple's LaserWriter 8 driver.

If you need to **install additional driver software** for a printer not made by Apple, follow the procedure described earlier in this chapter under "Install a Mac Printer Driver."



Using Shared Printers from a Macintosh

Select a Printer Driver



For more information on printing on the Mac, consult the Mac OS on-screen help. It's available from the Help menu (the question-mark menu prior to Mac OS 8) when the Finder is the active application.

More About
Selecting Drivers

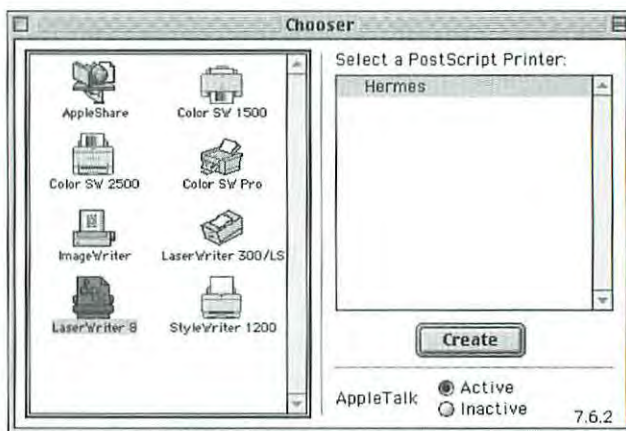


Select a Specific Printer

Attention StyleWriter and LaserWriter 300/LS Users:

To print on a shared local printer that uses any of the StyleWriter drivers or the LaserWriter 300/LS driver, the Mac that's going to print must have exactly the same driver version as the Mac with the shared local printer. You can check the driver version by selecting the icon of the printer driver in the Extensions folder and then choosing the Get Info command from the File menu. If necessary, copy the newest driver onto a floppy disk and use it to replace older drivers on other Macs.

4. After selecting a printer driver on the left side of the Chooser, you select a specific printer on the right side of the Chooser. For independent network printers, you'll see a list of printer names. Shared local printers are listed by name along with the Mac's printer and modem ports.



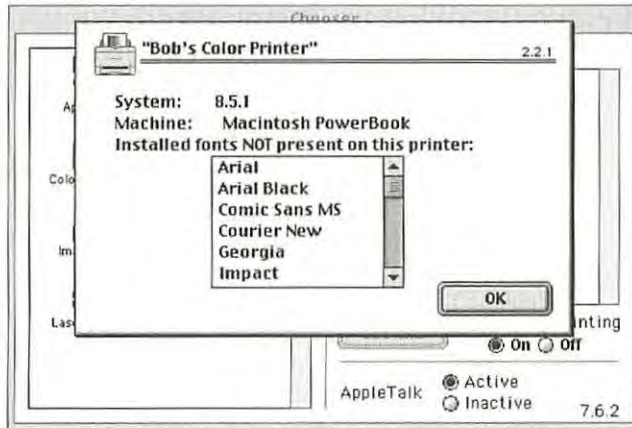
Click the LaserWriter 8 icon to see the names of PostScript printers available on the network.

Set Up a PostScript Printer

5. If you select a LaserWriter 8 printer in the Chooser, a Create button or a Setup button appears below the list of PostScript printer names. Do one of the following:
 - If the **Create** button appears, click the button to begin an automatic setup process. If you don't go through the setup process, the LaserWriter 8 driver uses generic settings based on the page size and features of the original LaserWriter printer.
 - If the **Setup** button appears, you can click it to change the printer's setup.

6. If you select a shared local printer (not an independent network printer) in the Chooser, a Get Info button appears below the list of ports and printer names. You can click this button to see the name of the selected printer's Mac connection, thus confirming that you have selected the printer you had in mind.

You'll also see a list of fonts that are installed on the Mac you're using but that are not present on the printer's Mac. (If you print a document containing fonts that are not present on the printer's Mac, it may print slowly or incorrectly.)



Click Get Info in the Chooser to see the name and other facts about a shared local printer's Mac.

7. When you finish selecting a driver, selecting a specific printer, and doing any necessary printer setup, you should close the Chooser. The printer you selected is now the default printer and will be used by Page Setup and Print commands in all Mac applications.

For some printers, an icon appears on the desktop after you select the printer in the Chooser. This **desktop printer icon** normally appears for printers that use any of the drivers included with the Mac OS (but usually not for printers that use other drivers). A heavy black border around the icon indicates the default printer.

You can use desktop printer icons to change the default printer and take care of other printing chores. You can also change the default printer with the Control Strip, if the Control Strip is available on your Mac.

Get Local Printer Info

Close the Chooser

Desktop Printer Icon



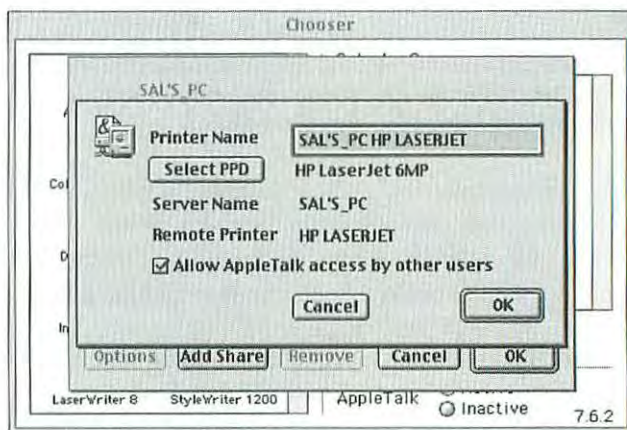
Connecting Macs to Shared PC Printers

If your network includes PCs as well as Macs, it's not normally possible for the Macs to use a shared local printer that's connected to a PC. This is partly because Windows 95 and 98 do not normally use the same network protocol used by the Mac OS for printing. Furthermore, Macs don't include driver software for printers that connect to PCs (with the exception of PostScript printers). You can overcome these obstacles by installing additional software on your Macs.

Use DAVE

Thursby Systems's **DAVE** software enables a Mac to use PostScript printers and files shared by PCs on the same network. DAVE is a bit complicated to set up, and we don't recommend it as a solution for just printing. But if you installed it for file sharing with PCs (as described in Chapter 7), then it'll also let your Macs use PostScript printers that are otherwise available only to the PCs on your network.

DAVE implements Microsoft networking software protocols on a Macintosh. With DAVE installed, PostScript printers connected to and shared by PCs will appear in the Chooser when you select the LaserWriter 8 icon. However, DAVE does not enable Macs to print to non-PostScript printers shared by PCs. It doesn't supply any printer drivers.



DAVE software enables Macs to use PostScript printers shared by PCs.

Another option is to install Miramar Systems's **PC MACLAN** on the PC. The PC MACLAN Print Server program lets Macs use the same printers used by PCs on the network, provided the Macs have appropriate driver software.

PostScript printers made available this way can use the LaserWriter 8 driver. These printers simply show up in the Chooser when you select the LaserWriter 8 icon. But PC MACLAN by itself doesn't enable printing to non-PostScript printers shared by PCs. For that, you need to install additional driver software on the Mac. (PC MACLAN also lets Macs and PC share files, as described in Chapter 7.)

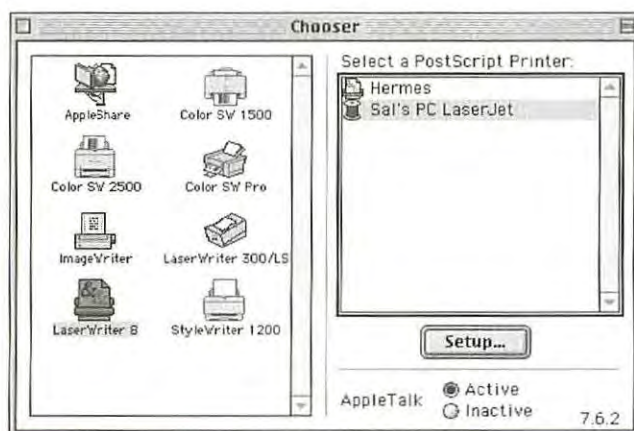
Use PC MACLAN



Sal's PC LaserJet



PC MACLAN's Print Server makes PostScript printers connected to PCs available on Macs.



The Chooser shows PostScript printers that are connected to PCs and shared by PC MACLAN.

Non-PostScript PC Printers

Ordinarily, a Mac can't print to shared printers connected to PCs if they don't use the PostScript printer language. A Macintosh simply doesn't have the drivers needed to use these printers.

You can, however, add printer drivers with a product called PowerPrint Networks from InfoWave (<http://www.infowave.com>). It supplies a Mac with printer drivers for more than 1600 printers that use the PCL printer language and other non-PostScript printer languages.


Summary

The desire to share a printer among several computers is reason enough to install a network. A network gives everyone access to a top-notch printer yet costs less than equipping every computer with a mediocre printer.

- ▼ You can share **local printers** that are connected directly to computers and independent **network printers** that are connected directly to your network.
- ▼ To share a printer, you'll need a physical connection for it as well as software for the network connection and the **printer driver** software.
- ▼ **To share a local printer with other PCs**, you need to plug the printer into the PC's parallel port or USB port; install the driver software on the PC; install the Windows network software for file and printer sharing; set parameters for Windows control of access to your shared printer; and designate the local printer as a shared printer.
- ▼ **To share a local printer with other Macs**, you need to plug the printer into a Mac serial port; install the driver software on the Mac; and designate the local printer as a shared printer.
- ▼ **Network printers** are easier to set up than local printers, because they're inherently shared devices. You simply plug the printer into your network and perhaps assign it an IP address. Then it's available to the computers on your network.
- ▼ Before a PC can use a shared local printer or a network printer that appears in its Network Neighborhood, the **printer icon must appear in the PC's Printers folder**. You can add a shared printer from the Printers folder or the Network Neighborhood in Windows 95 or 98. You can also enable PCs to use AppleTalk printers by installing additional software.
- ▼ To use a **shared printer on a Mac**, you use the Chooser to designate the shared printer as the default printer. (The Mac must have the compatible printer driver in its Extensions folder.) You can install software that enables Macs to use some printers shared by PCs.

Share Files with Others

A network makes floppy disks all but obsolete. No longer must you walk floppies or Zip disks from one computer to another to share your files. Instead you can set up **file sharing** on your network—which lets you copy, open, and save files on other computers as easily as you can on your own hard disk. Don't worry; setting up file sharing doesn't necessarily mean that it's open season on your files and everyone using your network has access to them. You can restrict access to your data in several ways. For example, you can allow access only to specific folders on your computer rather than to your entire hard disk, or you can require that a password be entered to gain access to your computer's hard disk.



Although the procedure for sharing files is somewhat different between a Windows PC and a Macintosh, you can share files across platforms. This chapter explains how to make files on one computer available to the other computers on your network. The other half of file sharing—accessing those shared files from another computer—is covered in Chapter 8.

Sharing Files with Other Windows PCs

Both Windows 95 and 98 include all the software a PC needs to share files with other Windows PCs on your network. Using this software involves completing the following tasks:

- Make sure the network software for file sharing is installed on the computer.
- Set the parameters for how Windows will control access to all the computer's shared folders and disks.
- Designate those folders or disks whose contents you want to share.
- Set access restrictions for each shared folder or disk.
- Monitor file-sharing activity.

The first four of these tasks must be done on each PC whose files you want to make available on your network; the last task is optional. None of these tasks is necessary for a PC that will *not* make its own files available but will use shared files of *other* computers.



Before you set up file sharing, make sure the network adapter, protocol, and client software are installed and configured, as described in Chapter 3.

Installing the File Sharing Service

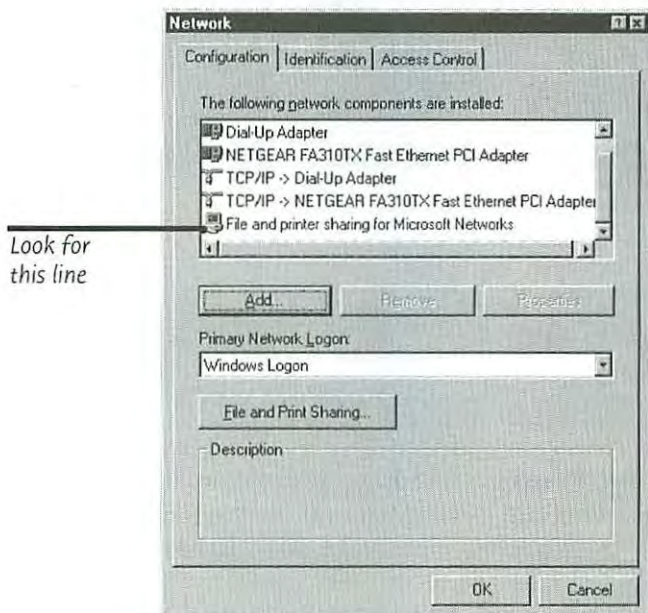
The **Windows file-sharing software** may already be installed on the PC. For instance, a PC that's set up to share its local printer (as described in Chapter 6) already has the necessary software installed. It's also possible that the file sharing software was installed when Windows was installed. In either case, you need to make sure file sharing is configured correctly. Use the **Network dialog box** to see if the PC has been properly set up with file sharing software.

The Windows software that allows a PC to share its local printer also allows the PC to share its files.

To install Microsoft File and Print Services:

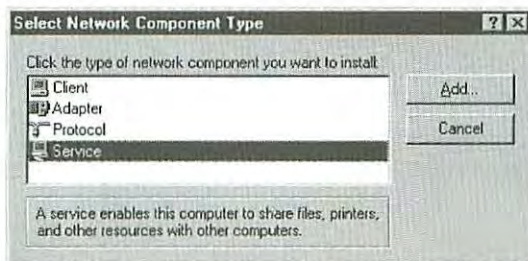
1. Open the Control Panel and double-click the Network icon. The Network dialog appears with the Configuration tab on top.

2. Look for "File and printer sharing for Microsoft Networks" in the list of installed network components. If you find this network component listed, skip ahead to the section "Setting Access Control." If you don't see this network component listed, you need to install the network component that provides the file sharing service.



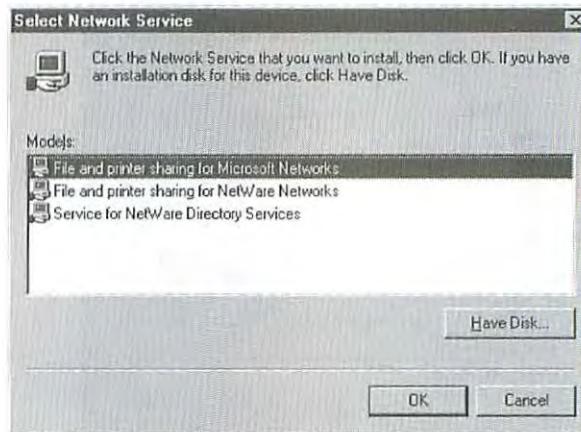
Check to see whether file and printer sharing software is already installed.

3. Click Add to open the Select Network Component Type dialog box.
4. Choose Service as the type of network component you want to install, and then click Add.

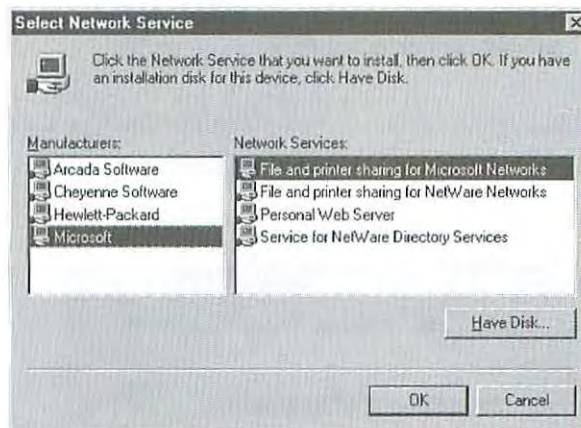


Choose Service as the type of network component to add.


5. In the next dialog box, you'll select the File and Printer Sharing software, as follows:
- **Windows 98 users** see a dialog box listing several services. Select "File and printer sharing for Microsoft Networks" and click OK.



- **Windows 95 users** see a dialog box with two lists. In the Manufacturers list, select Microsoft; then select "File and printer sharing for Microsoft Networks" in the Network Services list. Click OK.



6. Windows may now ask you to insert your installation CD-ROM so it can copy the file sharing software to your hard disk. Go ahead and follow the prompt.
7. Scroll the list at the top of the Configuration tab of the Network dialog box. You should see "File and printer sharing for Microsoft Networks" listed.
8. Click the button labeled File and Print Sharing.
9. In the File and Print Sharing dialog box, select the option "I want to be able to give others access to my files." (For file sharing, it doesn't matter whether the other check box is selected or not.) Click OK to return to the Network dialog box.

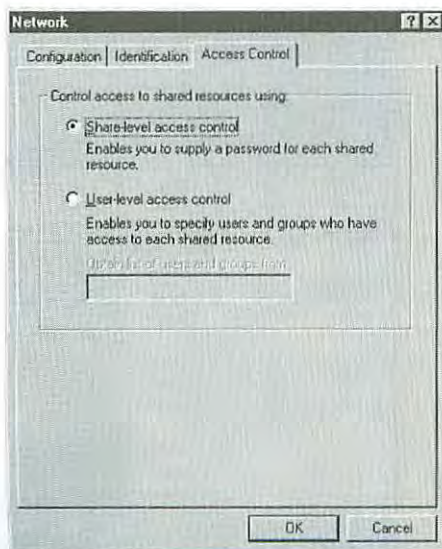


The next step in making your PC's files available for sharing is to determine how Windows controls access to the shared folders. On a peer-to-peer network, you can control access only by assigning passwords to each shared folder and disk, as described shortly. This is called **share-level access control**.

Setting Access Control

To set the access control method:

1. In the Network dialog box, click the Access Control tab to open it.
2. Select the option for share-level access control.



Share-level access control is the only method available on a network without a Windows NT server.

- Click the OK button to close the Network dialog box.

Windows may ask you to insert the Windows CD so that some files can be copied to your hard disk. After copying, Windows will ask you to restart your computer. You can do that now, or press on.



The other method of controlling access to shared items (**user-level access control**) is not available unless your network has a Windows NT server. Server-based networks are not covered in this book. If you add a Windows NT server to your network, you can select user-level control in the Access Control tab of the Network dialog box. In this case, each person who uses the network has an individual password, and you can specify which users or groups of users have access to each shared folder and disk.

Selecting Folders to Share

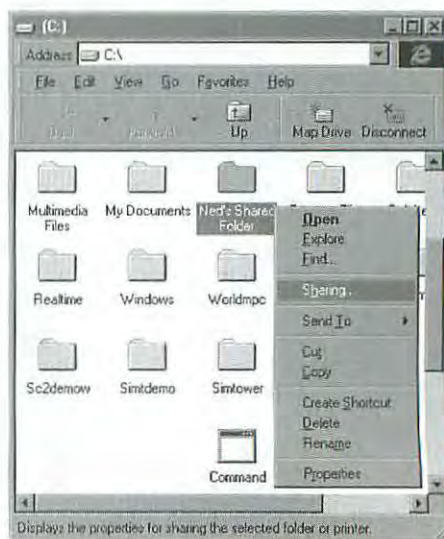


Windows uses the term “share” to describe any directory (that is, any folder or disk) that you allow other computers to access over your network. The process of making a folder or disk a share is called **creating a share**. (The same term applies to shared printers, as described in Chapter 6.)

Before making a folder a shared item, think carefully about its name and location. If you move or rename the folder, or any folder or disk in its path, the shared item will no longer be shared. Avoid changing a folder’s path after making the folder a shared item.

To share a folder or disk:

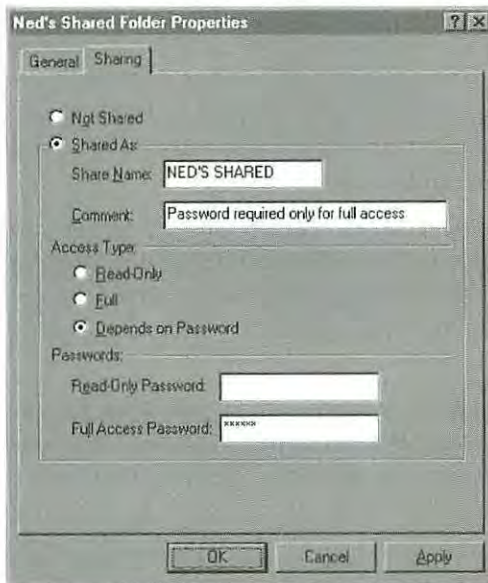
- Locate the icon of the folder or disk you want to share, either in Windows Explorer or by clicking My Computer and finding it there.
- Right-click the icon to open its shortcut menu. Then select Sharing.
- Click the Shared As option on the Sharing tab in the Properties dialog box.
- In the Share Name field, type in a name for the folder or disk. This name



will appear in the Network Neighborhood of all PCs on your network. You can use the same name as the folder or disk itself, or you can call it something different. The name can have no more than 12 characters.

When naming a folder or disk, you can hide the item from the Network Neighborhood of other (unauthorized) users by appending a dollar sign (\$) to the name. This gives you greater control over who can use the shared item, because anyone who wants to access it will have to know its exact name and will need to map a drive letter to it (see Chapter 8 for more about this).

5. In the Comment field, type a phrase that describes the shared item or its location. If this PC has more than one shared folder or disk, the comment will help users identify this one.



The Share Name and Comment can be seen by other network users; passwords are optional.

Changing the name of a share on your computer can make it more difficult for others to use. That's because networked PCs will still display the old share name in the drop-down Path list of the Map Network Drive dialog box, because they don't know that you changed the name (as described in Chapter 8.) If someone tries to select the old name from the share list, they'll get the error message "Share name not found." To use the renamed shared item, they'll have to learn the new name and type it in.



Restrict Access to Files*Look for the Hand Icon*

6. Choose an option to restrict access to the share. The bottom half of the Sharing tab in the Properties dialog box contains options for **restricting access**. Under the heading Access Type, you'll see three choices that restrict access: Read-Only, Full, or Depends on Password.


Cable Modem Users: Always specify passwords if you use a cable modem connection to the Internet. If you leave a password field blank, any neighbors with PCs and cable modems will be able to access your shared folders on their PCs.

- **Read-Only:** Users on other network PCs will be able to see and open the files in a Read-Only shared folder or disk, but they won't be able to make changes to the files or copy new files into the shared folder or disk. Nor will they be able to delete files from the shared folder or disk. They will be able to copy files from the shared folder or disk to their own hard disks.
- **Full:** This option allows people using other network PCs to read and write to files on this shared folder or disk. In other words, network users can see and open, copy, save changes to, and delete files on this folder or disk.
- **Depends on Password:** This access type lets you set two different passwords for connecting to the shared folder or disk—one password is for Read-Only access, and the other grants Full access. Each password can be no more than eight characters in length.

To allow Read-Only access without a password and Full access with a password, set the access type to Full, leave the Read-Only Password blank, and then enter a password for the Full access Password.

7. Click OK to finish making the folder or disk a shared item. Take a look at the shared item's icon, and you'll see that it now sports a hand image to indicate the item is shared. The folders and disks you designate as shared items will appear in the Network Neighborhood of the PCs on the network. (Chapter 8 explains how to connect to shared folders and disks over the network.)

You can see who is currently connected to a PC's shared items by opening the **Net Watcher utility**. This utility is normally buried several levels deep in the Start menu (point to Programs, then Accessories, System Tools, and then Net Watcher).



Oversee shared items with Net Watcher.

Net Watcher shows you what shared items are available on the PC, who is currently connected to them, and which files are open on other network computers. You can change your view of the network activity by clicking one of the buttons in the toolbar or by selecting an equivalent command in Net Watcher's View menu.

Other toolbar buttons and menu items allow you to take action on connected users, shared folders, and open files. As detailed in the following paragraphs, these controls allow you to disconnect individual network users, close a file that someone has open, stop sharing any shared item, share additional items, and change access restrictions, right from your computer.

Click the **Show Users button** to list the users who are currently accessing this PC's shared items. From this list, you can select a user to see which shared items the user is connected to.

To disconnect a user, select the user in the list and then click the **Disconnect User** button on the toolbar.

Click the **Show Shared Folders button** to see a list of all the shared items that are available on this PC. From the list, you can select a shared item to see whether anyone is connected to it.

You can perform several operations on the shared folders listed here. To stop sharing an item, select it and click the **Stop Sharing button** on the toolbar. To change access restrictions from this list, select the item you want to change and then select **Shared Folder Properties** from the Administer menu. To make another shared item, click the **Add Share button** on the toolbar and specify the path of the item you want to share.

Monitoring Shared Items

Show/Disconnect Users



Show Shared Folders



Show Files

Click the **Show Files button** to list this PC's files (from its shared folders) that are now open on another network computer. You can close a file by selecting it in the list and then clicking the **Close File button** on the toolbar.

Sharing Files with Other Macs

A Macintosh includes all the software it needs to share files with other Macs on your network. **File sharing software** is a standard part of the Mac OS (since System 7, circa 1991), so you won't need to install anything unless you have a really old Mac. Of course, you do need to set up the file sharing software. On each Mac whose files you want to share, you need to perform the following simple tasks:

- Turn on file sharing.
- Designate the folders and disks whose contents you want to let others share.
- Set access privileges for each shared folder and disk.
- Set the Guest access for the whole computer.

These tasks are necessary only for Macs whose files will be shared. None of these tasks is necessary on a Mac that will not share its own files but will use shared files from other computers. Let's take a look at these Mac file-sharing tasks.



You don't need to restart a Mac after you change its file sharing setup. This includes turning file sharing on and off, making a folder or disk shared or not shared, and changing a shared item's access privileges.

Turn File Sharing On/Off

Until file sharing is turned on, a Mac can't share its files with other computers on the network. To turn on file sharing, go to the File Sharing control panel of Mac OS 8 or later, or the Sharing Setup control panel of Mac OS 7.6.1 or earlier.

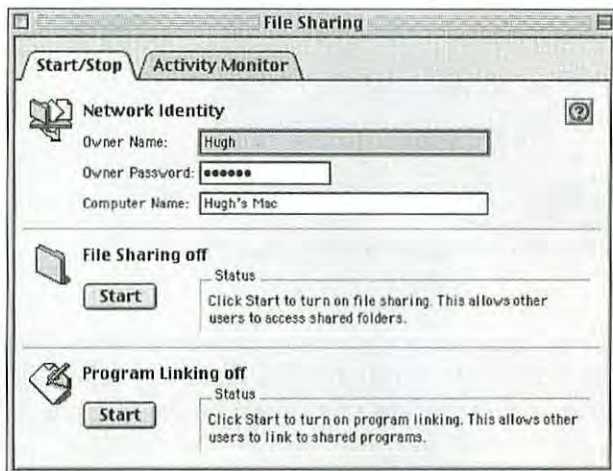
Other Shared Services:

Turning file sharing on or off has no effect on other network services, such as sharing an Internet connection or playing games.

To turn on Mac OS file sharing:

1. Open the File Sharing control panel. Click the Start/Stop tab (Mac OS 8 or later), or open the Sharing Setup control panel (Mac OS 7.6.1 or earlier). Make sure the Owner Name, Owner Password, and Computer Name fields are filled in (as described in Chapter 4).
2. Click the upper Start button (the one under "File Sharing Off" or "File Sharing"). The button changes to a Cancel button while file sharing starts up, which may take a little while. When file sharing is on, the button becomes a Stop button.

Program Linking: You can ignore the Program Linking section at the bottom of the control panel. It has no effect on file sharing.



The File Sharing control panel in Mac OS 8 and later.



The Sharing Setup control panel in Mac OS 7.6.1 and earlier.

When you turn off your Mac's file sharing capability, you prevent other computers from using its shared items, but your computer can still use shared items on other computers.

To turn off file sharing, click the Stop button in the File Sharing or Sharing Setup control panel. A dialog box appears, in which you set a number of minutes until file sharing is turned off. During this interval, your computer notifies other networked computers about the change. Users on computers that are connected to your computer's shared items are notified of how long they have until the "file server" shuts down. You can set the delay period to zero to stop file sharing immediately, but users on other computers will be disconnected without warning.

Turning Off File Sharing

You can set a polite interval before file sharing turns off.



If your Mac has a **Control Strip**, you can use it to turn file sharing on and off. Just click the File Sharing button, which looks like a folder icon, and choose the appropriate command from the pop-up menu.



The File Sharing button in the Control Strip.

Selecting Folders to Share



Once Mac file sharing is turned on, you can designate shared items. You can share folders and disks, including hard disks, CD-ROMs, Zip disks, and most other removable disks that appear on the Mac's desktop. However, you can't share floppy disks or the folders on them.

You can designate up to 10 folders and disks as shared items. Folders nested inside a shared item do not count toward this limit. Network users can connect to the shared items via the Chooser or Network Browser, as described in Chapter 8.

Nested Folders: You cannot make a folder a shared item if it is nested within a folder that is already a shared item. Network users can access the inner folder by connecting to the outer shared folder, or you can move the inner folder to the desktop or to another folder, where it can be made a separate shared item.

To share a folder or disk:

1. Select the folder or disk icon in the Finder.
2. Display the sharing info for the selected folder or disk:
 - **Mac OS 8.5 and later:** From the File menu, choose Get Info and then Sharing.
 - **Mac OS 8.1 and earlier:** From the File menu, choose Sharing.



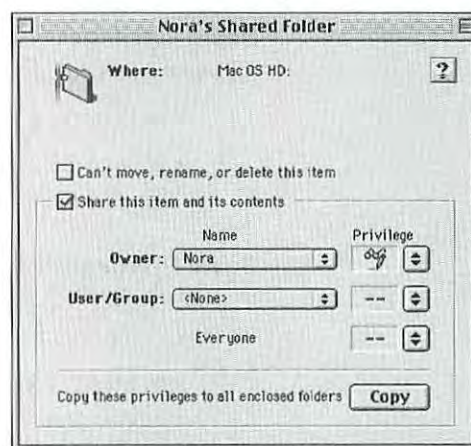
In Mac OS 8 and later, you can use a **contextual menu** to display the sharing info for a folder or disk. Simply Control-click the folder or disk icon to pop up its contextual menu. Then choose Sharing from the Get Info submenu (Mac OS 8.5 and later) or choose Sharing directly from the contextual menu (Mac OS 8–8.1).

3. Select the option "Share this item and its contents."

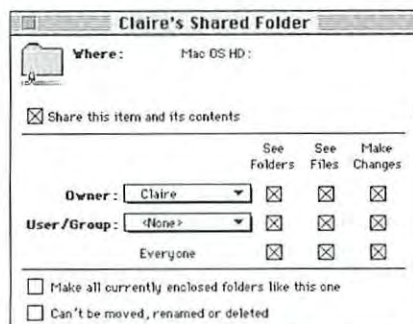
- If you want to lock the shared folder, select the option "Can't move, delete, or rename this item." A locked folder can't be moved to another folder, be dragged to the Trash, or have its name changed by anyone. Enabling this option does not, however, prevent anyone from dragging files and folders into or out of the shared folder.



A folder's sharing info in Mac OS 8.5 and later.



A folder's sharing info in Mac OS 8-8.1.



A folder's sharing info in Mac OS 7.6.1 and earlier.

You can lock any folder by turning on the option "Can't move, delete, or rename this item" in its sharing info. The folder does not have to be shared (the option "Share this item and its contents" does not have to be selected) to be locked.



Setting Basic Access Privileges

To limit what network users can do with the contents of a shared folder or disk, you can set **access privileges**. You have a great deal of flexibility in setting access privileges, as described in the upcoming section “The Ins and Outs of Access Privileges.” But if you’d rather keep it simple and avoid the details, you can set privileges according to one of the following basic scenarios:

Basic Privileges for Minimum Security: This arrangement gives all network users full access to the shared item’s contents:

- In Mac OS 8 and later, set Everyone to Read & Write.
- In Mac OS 7.6.1 and earlier, turn on all the Everyone privileges, including See Folders, See Files, and Make Changes.

Basic Privileges for Maximum Security: In this arrangement, you require network users to know a Mac’s owner name and password to use the contents of a shared folder or disk.

- In Mac OS 8 and later, set Everyone to None, User/Group to None, and Owner to Read & Write.
- In Mac OS 7.6.1 and earlier, turn on all of the Owner privileges and turn off all the User/Group and Everyone privileges.

If you use the maximum security setup described here, you can make it easier for network users to use shared items by entering the same owner name and password in the File Sharing or Sharing Setup control panel on all Macs in the network. (Each Mac still must have a unique computer name, however.) Having a common owner name and password isn’t as secure as having different names and passwords, but it’s a lot less stuff to remember.

Once you’ve shared a folder, its icon sprouts a network wire to let you know it is shared. In addition, faces appear on a shared folder’s icon when someone is connected to it from another Mac. The icon for a disk, however, does not change when you make it a shared item.

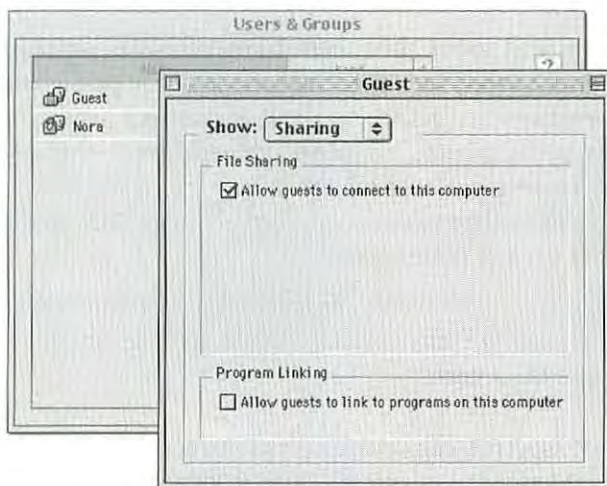


Setting Guest Access

In addition to setting access privileges for each shared folder and disk, you need to determine whether to allow guests to connect at all. **Guests** are unidentified network users who can connect without a password. In Mac OS 8 and later, by default, guests are *not* allowed to connect; but in Mac OS 7.6.1 and earlier, they are allowed to connect.

To turn Guest access on or off:

1. Open the Users & Groups control panel.
2. Double-click the Guest icon in the Users & Groups window.
3. From the Show pop-up menu in the Guest window, choose Sharing. (In Mac OS 8.1 and earlier, you won't need to do this step.)
4. If you want to allow unidentified network users to connect, turn on the option to "Allow guests to connect to this computer"; otherwise, turn it off.
5. Close the Guest window and the Users & Groups window.



Turn Guest access on to allow any network user to use shared folders and disks without a password.

Attention Cable Modem Users: If you have a cable modem connection to the Internet, make sure Guest access is **turned off**. Remember that every Mac in your neighborhood with a cable modem is part of your local network. With Guest access turned off, your neighbors can't connect to a shared folder or disk unless they know a name and password that's listed in your Mac's Users & Groups control panel. This warning does not apply to other Internet connection methods, including regular modems, ISDN, and DSL.





Protection for Shared Files: If you'd like to simplify your life by turning on Guest access but are concerned about the security of some of your Macs, keep in mind that you can protect a Mac's data by turning off file sharing. For instance, you can enable file sharing on the Macs in your kids' rooms but disable it on your home office Mac. This lets you move files between your Mac and the shared folders on your kids' Macs, but it prevents your kids from getting into folders or disks on your Mac. And keep in mind that you don't have to have file sharing turned on all the time. You can turn file sharing on, move some files, and then turn it off.

The Fine Points of Mac File Sharing

What you've read so far about controlling access to shared items on your Macintosh barely scratches the surface of the possibilities. Many other alternatives are available to you for setting access privileges, adding users and groups, monitoring shared items, inheriting access privileges, keeping enclosed folders private, changing ownership of a folder, and getting the same users and groups on all Macs. In this section we'll explore each of these topics.

If you're happy with the simple access control alternatives described so far in this chapter, you can skip ahead to "File Sharing Between Windows and Macintosh."

Ins and Outs of Access Privileges

If you're not satisfied with setting access privileges according to one of the basic scenarios, or you simply want to know more about how privileges work, you'll need to learn about the categories of users and the privilege levels you can set for each category. Then you'll be able to get more creative in limiting what network users can and can't do with a shared item.

User Categories

As explained in the following sections, you can set explicit access privileges for three categories of users: Owner, User/Group, and Everyone.

- **Owner:** Initially, the owner of a folder or disk is the person named as the owner in the File Sharing or Sharing Setup control panel. This is usually the person who works on the Mac most of the time. If the Mac is used by several people, the owner could be the person who manages your network (probably you). If you specify the same owner name and password on all your Macs, everyone who knows the password will have owner privileges.

- **User/Group:** You can designate one network user or a group of users to get special access privileges. To utilize this category, you must register users and groups in the Users & Groups control panel, as described shortly in “Adding Users” and “Adding Groups.”
- **Everyone:** The Everyone category includes users who are registered in the Users & Groups control panel and who may have passwords. This category also includes guests (unidentified users without passwords) if Guest access is enabled in the Users & Groups control panel.


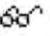


You don't have to set privileges for the Owner or User/Group categories if they have the same privileges set in the Everyone category, especially if the Everyone category has full access privileges.



For each of the user categories, you can set explicit access privileges. In **Mac OS 8 and later**, you choose a privilege level from a pop-up menu. In **Mac OS 7.6.1 and earlier**, you set privileges by turning check boxes on and off in any combination. The following two tables explain the access privilege settings according to their respective Mac version.

Privilege Levels

Access Privilege Levels (Mac OS 8 and Later)

Privilege Level	Permissions Granted
 Read & Write	Users can see and open all the files and folders in the shared folder or disk and can copy files from it. Users can also save changes to files and add files.
 Read only	Users can see and open all the files and folders in the shared folder or disk and can copy files from it. Users can't save any changes to files or copy files to the shared folder or disk.
 Write only	Users can add files to a shared folder or disk but can't open it.
 None	Users can't see, open, or add to the shared folder or disk.

Access Privilege Settings (Mac OS 7.6.1 and Earlier)

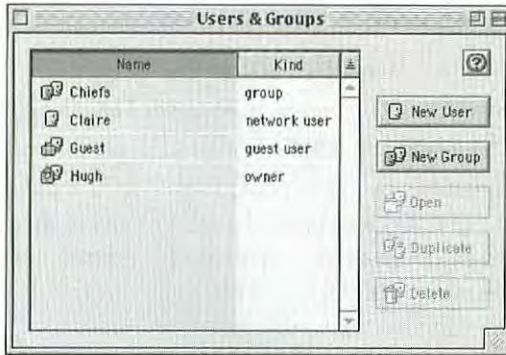
Privilege	Permissions Granted	Similar Privilege in Mac OS 8 and Later
See Folders	Users can open the shared folder or disk, can see and open folders in it, and can copy folders from it.	N/A
See Files	Users can see and open files in the shared folder or disk and can copy files from it.	N/A
Make Changes	Users can add files and folders to the shared folder or disk. Users can't open the shared item to create, delete, move, and change files or folders unless See Files or See Folders is turned on.	Write only
See Folders and See Files	Users can see and open files and folders in the shared folder or disk and can copy items from it.	Read only
See Folder, See Files, and Make Changes	Users can add files and folders to the shared folder or disk; and can create, delete, move, and change files and folders that are inside the shared folder or disk.	Read & Write

When You Upgrade...


Several possible combinations of privilege settings are available in Mac OS 7.6.1 and earlier that don't exactly correspond to privilege levels in Mac OS 8 and later. However, Mac OS 8 and later will honor settings made in versions 7.6.1 and earlier, allowing your network to include Macs with older and newer versions of the Mac OS. If Mac OS 8 or later encounters a combination of privilege settings that don't correspond to one of its privilege levels, the system will display a question mark icon instead of a glasses and pencil, glasses only, or pencil only icon.

Adding/Removing Users

As described earlier, you can allow everyone to access a Mac's shared folders and disks, or you can allow only the owner to have access. But you don't have to settle for only these two alternatives. You can use the **Users & Groups control panel** to identify individual network users and groups of network users, granting or denying them special access privileges to your shared items.



The Users & Groups control panel in Mac OS 8 and later.



The Users & Groups control panel in Mac OS 7.6.1 and earlier.

Each user is represented in the control panel by an icon, and each user except Guest has a unique name and can have a password. The Owner name and password are taken from the File Sharing or Sharing Setup control panel. Initially, each Mac has two users in its Users & Groups control panel: the Owner and Guest. When you add new network users to the Users & Groups control panel, you assign their names and passwords as well.

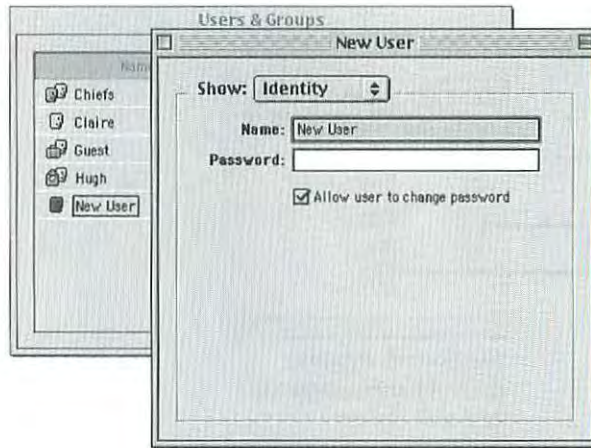
Macintosh file sharing, unlike its Windows counterpart, doesn't need a dedicated server to have **user-level access control**. Each Mac stores its own list of users who will access it.



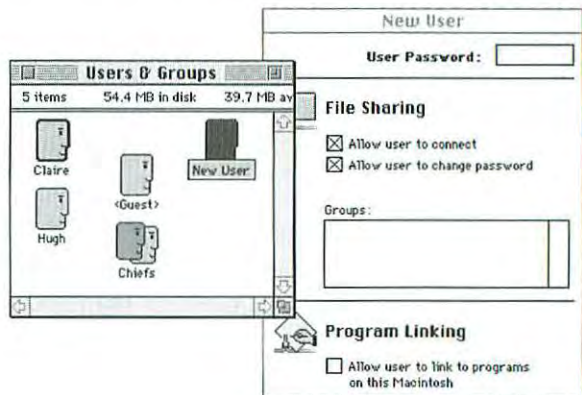
To add a network user:

1. Open the Users & Groups control panel.
2. Click the New User button (or, in Mac OS 7.6.1 and earlier, choose New User from the File menu).

An icon named New User appears in the control panel, and if you're using Mac OS 8 or later, the icon opens to display the user's window and the user name is selected for editing. If you're using an earlier version, the icon does not open but the user name is selected for editing below the icon.



Adding a new Macintosh file sharing user in Mac OS 8 and later.



Adding a new user in Mac OS 7.6.1 and earlier.

3. Type a name for the new user. You can use any name, but it's most convenient to use the same name that this person has entered in the Owner Name field of his or her own File Sharing or Sharing Setup control panel. This is the name that will be supplied automatically when the user tries to connect to another Mac, as described in Chapter 8.

For simplicity and clarity, make sure each registered user has the same name on all network Macs. Don't use "Elizabeth" on one Mac, "Liz" on another, and "Beth" on a third, for example.



4. If the user's window is not already displayed, open the user's icon. Type a password in the user's window; use a different password for each person. You can leave the password blank if you don't want to require the user to enter one.
5. If you want, turn on the option to "Allow user to change password"; the user will be able to replace the password with one that's easy for him or her to remember.
6. Close the user's window.

User names and passwords are meant to identify individuals, not computers. If several people share the same computer, add a separate user item in the Users & Groups control panel for each person. But keep in mind that a user name and password are not positive proof of identity. Anyone who learns someone else's user name and password can access your shared folders and disks as an impostor.



To remove a network user:

- ▼ Drag the user's icon to the Trash.
or
- ▼ Select one or more user icons and click the Delete button or choose Delete from the File menu. (The Delete button does not exist in Mac OS 7.6.1 and earlier.) Users are deleted immediately. They don't wait in the Trash until you empty it, as do many other deleted items.

Adding Groups

With Macintosh file sharing, you can add groups of network users and give the group special access privileges for a shared folder. For instance, you could give a group full access to a folder to which everyone else has partial access.

To add a group:

1. Open the Users & Groups control panel.
2. Click the New Group button (or, in Mac OS 7.6.1 and earlier, choose New Group from the File menu).

An icon named New Group appears in the control panel, and if you're using Mac OS 8 or later, the icon opens to display the group's window with the group name selected for editing. If you're using an earlier version, the icon does not open but the group's name is selected for editing below the icon.

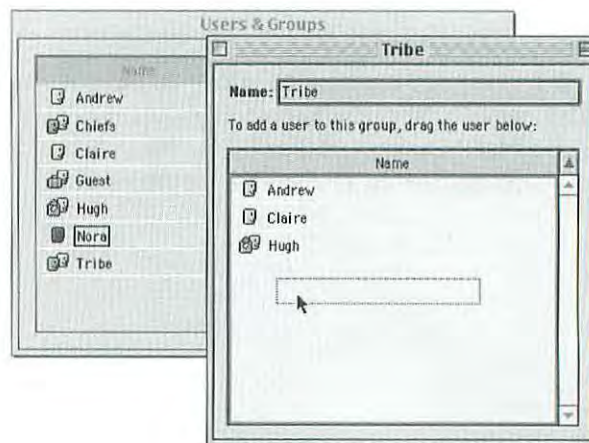
3. Name the group.

Adding Users to a Group

You can add members to a group in any of these ways:

- Drag users from the Users & Groups window into the group window.
- Drag users to the group icon in the Users & Groups window.
- In Mac OS 8 and later, drag one group to another group, in which case all the members of the dragged group become members of the other group.

You can select several users (as you would select multiple icons in the Finder) and drag them together to a group icon or window.



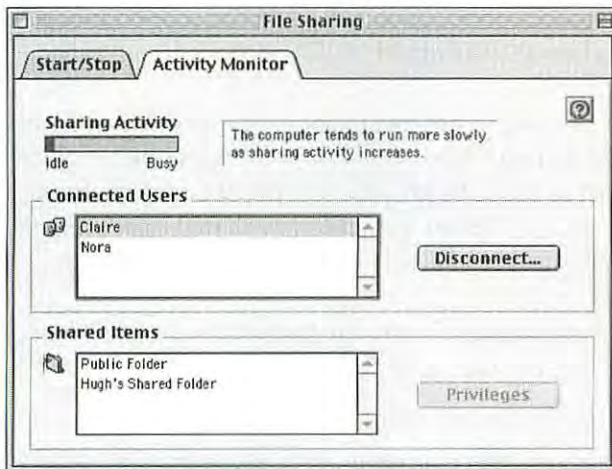
Add members to a group by dragging their icons to the group window.

To remove a user from a group, drag the user's icon from the group window to the Trash. If you're using Mac OS 8 or later, you can also remove users by selecting their icons in the group window and pressing the delete key, or by choosing Remove from the File menu.

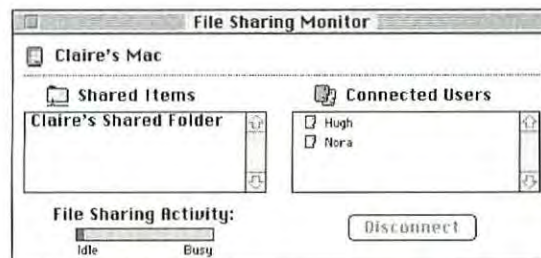
Once you set up file sharing, it hums along in the background. Network users connect to your Mac's shared folders without your knowledge. If you're ever wondering who is connected, or you want to review which folders and disks are shared items, you can use the Activity Monitor tab of the File Sharing control panel (if you use Mac OS 8 or later). In Mac OS 7.6.1 and earlier, a separate control panel called File Sharing Monitor provides the same information.

Removing users from a group

Monitoring File-Sharing Activity



In Mac OS 8 and later, File Sharing's Activity Monitor tab shows which items are shared and who is connected.



In Mac OS 7.6.1 and earlier, the File Sharing Monitor shows which items are shared and who is connected.

Disconnecting Users

You can disconnect any user by selecting the user's name in the control panel's list of connected users and then clicking the Disconnect button. A dialog box asks how many minutes' delay you want to allow before the user is disconnected, and the user will be notified in advance of the approaching disconnection. You can enter 0 minutes to disconnect a user without warning.

Changing Privileges or Share Status

With Mac OS 8 and later, File Sharing's Activity Monitor tab makes it easy to change access privileges for, or even stop sharing, any shared item. Simply double-click the item in the list of shared items, or select the item and click the Privileges button, to display the item's sharing info (shown earlier in "Selecting Folders to Share"). There you can set different access privileges or stop sharing the item by turning off the option "Share this item and its contents." Close the sharing info window to make any changes take effect.

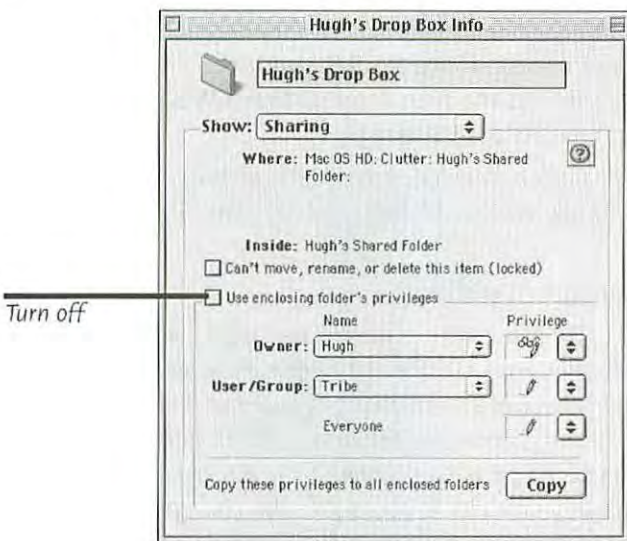
Inherited Access Privileges

Initially, all folders enclosed in a shared folder or disk inherit their access privileges from the shared folder or disk. If you change the privileges of the shared folder or disk, the privileges of the enclosed folders also change. Move an enclosed folder to a different shared folder, and the enclosed folder inherits the privileges of its new enclosing folder.

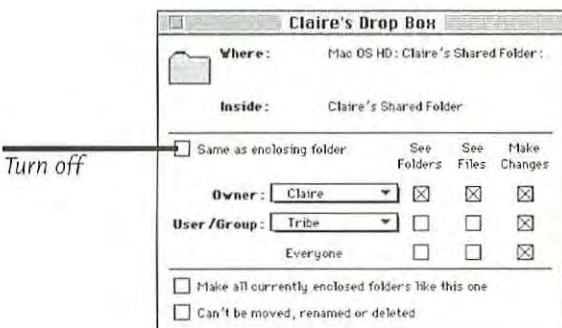
Setting Independent Privileges

You can explicitly set privileges for any enclosed folder, independent of the shared folder that encloses it, as follows:

1. Click to select the enclosed folder in the Finder.
2. Use the Sharing command to bring up the selected folder's sharing info. (**Mac OS 8.5 and later:** From the File menu, choose Get Info and then Sharing. **Mac OS 8.1 and earlier:** Choose Sharing from the File menu.)
3. Turn off the option "Use enclosing folder's privileges" (Mac OS 8 and later) or the equivalent option "Same as enclosing folder" (Mac OS 7.6.1 and earlier).
4. Set the privileges you want for the enclosed folder.
5. Close the sharing info window.



Setting independent privileges for a folder enclosed in a shared item (Mac OS 8 and later).



Setting independent privileges for a folder enclosed in a shared item (Mac OS 7.6.1 and earlier).

After you have set independent privileges for an enclosed folder, it retains them even if you move it to a different enclosing folder. To make an enclosed folder inherit the privileges of its enclosing folder, turn on "Use enclosing folder's privileges" or "Same as enclosing folder" (depending on the Mac OS version).

You can force folders to use the same privileges as the folder or disk that encloses them by taking these steps:

1. Select the folder or disk that has the privileges you want to force on its enclosed folders.

Forcing Inherited Privileges

2. Use the Sharing command to display the sharing info for the enclosing folder or disk. (**Mac OS 8.5 and later:** From the File menu, choose Get Info and then Sharing. **Mac OS 8.1 and earlier:** From the File menu, choose Sharing.)
3. Click the Copy button (Mac OS 8 and later) or turn on the option "Make all currently enclosed folders like this one" (Mac OS 7.6.1 and earlier).
4. Close the sharing info window.



When you force a folder to inherit privileges, this action coincidentally stops all affected folders from *automatically* inheriting any subsequent changes you might make to the privileges of the enclosing folder or disk. If you change the enclosing item's privileges, none of the enclosed folders is updated automatically. To propagate the change to the enclosed folders, you have to explicitly repeat the steps just above for forcing inherited privileges.

File Sharing Between Windows and Macintosh

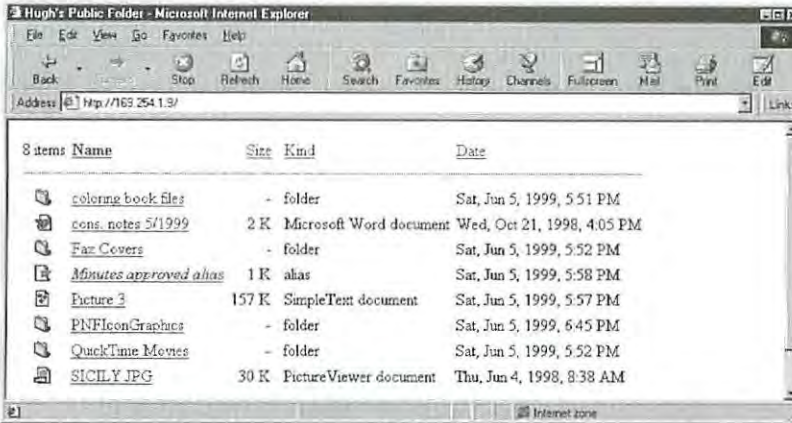
What we've covered so far about file sharing on Macintosh computers and Windows PCs won't help you set up file sharing between Macs and PCs on the same network. For that, you need more than the software and commands included with Windows and the Mac OS. You have three alternatives:

- Use Internet technology, such as a simple Web server.
- Add specific cross-platform file sharing software to the PCs or the Macs.
- Add a dedicated file server that can handle both platforms.

Let's take a closer look at each alternative.

Web Technology for Cross-Platform Sharing

Included with **Mac OS 8 and later** is a feature called **Personal Web Sharing** that can be used to share one folder or one disk with any other computer that has a Web browser. Obviously, this includes PCs with Windows. Personal Web Sharing creates a personal Web site on your network, and this site can list the contents of one Mac folder or disk. Network users can connect their Web browsers to this site to gain access to the contents of the folder or disk, including the contents of enclosed folders.



A Windows Web browser can access folders and files on a Macintosh that has the Personal Web Sharing utility. Notice that the Windows user sees file icons as they appear on the Mac.

Network users who access files through Personal Web Sharing will be able to open only certain types of files, including text files, within their Web browsers. For other types of files, such as Microsoft Word files, users will have to copy the files to their computers and then open the copies. In contrast, all other forms of file sharing allow users to open the original files directly.



Here's how to use Personal Web Sharing to share a folder:

1. Make sure the Mac has the TCP/IP protocol configured, as discussed in Chapter 4.
2. Turn on file sharing in the File Sharing control panel, as described earlier in this chapter.
3. Select the folder or disk you want to share. Then use the Sharing command (from the File menu, choose Get Info and then Sharing) to make it a shared item, as described earlier in this chapter.
4. Open the Web Sharing control panel.
5. Click the upper Select button, next to the words "Web Folder: ...". In the dialog box that appears, select the folder or disk you want to share.
6. Click the lower Select button, next to the words "Home Page: ...". In the dialog box that appears, click None. This setting indicates that you want Personal Web Sharing to list the contents of the folder you selected in step 5.

Setting Up Personal Web Sharing

Mac OS 8 and Later

7. At the bottom of the control panel, select the option "Use File Sharing to control user access." This tells Personal Web Sharing to control access to the folder or disk you selected in step 5 by using the access privileges you set up with ordinary Mac file sharing.
8. Click the Start button.



Set up the Web Sharing control panel to list the contents of one folder or file on your personal Web site.

Your Web Site's Address

After Web Sharing has started up, the address of your personal Web site appears at the top of the Web Sharing control panel in the My Address line. On most small networks, this address will contain a numeric IP address such as `http://169.254.1.9`, not a Web-site name such as `http://www.peachpit.com`. Tell network users your address so they can point their Web browsers to it and see the shared folder or disk and its contents.

Security Issues

User Name and Password: When network users try to connect to your personal Web site, they may be asked to enter a user name and password. This is the name and password you have set up in your Users & Groups control panel, as explained earlier in this chapter. Network users do not have to go through this authentication process, however, if you select the option "Give everyone read-only access" in the Web Sharing control panel.



Network users can be required to enter a user name and password to see the shared folder or disk on your personal Web site.

Connecting to the Internet with Web Sharing turned on opens the possibility of someone on the Internet accessing your shared folder or disk.

If all the Macs on your network share an Internet connection, it's possible that all the Macs with Web Sharing turned on will be open to intrusion by someone else on the Internet when any one of your networked Macs makes an Internet connection. However, if your Internet gateway has a firewall feature (as described in Chapter 5), you should be protected. If you're not sure about the firewall, you can prevent Internet invaders from accessing your Macs' personal Web sites by turning on the option "Use File Sharing to control user access" in the Web Sharing control panel. Then turn off Guest access in the Users & Groups control panel on each Mac.



The Personal Web Sharing solution to cross-platform file sharing is not as convenient as regular file sharing, and it doesn't let PC users share their files with Macintosh users. To enable your Macs to participate fully in Windows file sharing, like PCs, you can add some specific cross-platform file sharing software. Alternatively, you can add software to the PCs, enabling them to participate fully in Mac OS file sharing. Your choice will likely be based on whether you have more Macs than PCs, or vice versa. With either solution, the Macs and the Windows machines will be able to access each other's files.

Software for Cross-Platform Sharing

Thursby Software's DAVE makes a Mac capable of using Windows file sharing. A Mac with DAVE installed can connect to folders and disks shared by PCs on the network. Conversely, the Mac can create Windows-style shares that show up in the Network Neighborhood of

DAVE

PC MACLAN**Printer Sharing:**

DAVE and PC
MACLAN also work
for cross-platform
sharing of printers.
See Chapter 6.

**Adding a
Dedicated File
Server**

the PCs. Because you must buy DAVE for each Mac that needs cross-platform file sharing, this solution is a good way to add a couple of Macs to a mostly Windows network. For more information, contact Thursby Software at <http://www.thursby.com>.

Miramar Systems' PC MACLAN makes a Windows PC capable of using Macintosh file sharing. With PC MACLAN installed, a PC can connect to folders and disks shared by Macs on the network. In addition, the PC can designate Mac-style shared items that appear in the Chooser and Network Browser of the Macs. If you have fewer PCs than Macs, you'll buy and install less software by getting PC MACLAN for each PC that needs cross-platform file sharing. (You can contact Miramar Systems at <http://www.miramarsystems.com>.)

The third way to enable file sharing between Windows and Macintosh machines is to add a **file server** that serves both platforms. This solution is expensive, however, because you must buy the server software and install it on a dedicated server computer. No one works directly on the dedicated file server computer—it only stores and serves files for other computers on the network. The file server computer can be a PC or a Mac.

The most popular server software that runs on a PC is **Microsoft's Windows NT Server**. It includes software called Services for Macintosh, which you install separately from Windows NT itself. Services for Macintosh enables Macintosh users to access the same files as Windows users. If you don't have any experience or training with Windows NT, Services for Macintosh can be a bit complex, and you might want to consider hiring a consultant to set it up for you.

If you want a Macintosh-based server, **AppleShare IP 6.0** (and later) serves files to both Windows and Macintosh computers. This is a standard feature; you don't have to install any separate software to support Windows users. AppleShare IP is also the easiest server to set up. Where Windows NT Services for Macintosh treats Mac users separately, Windows and Mac users are served together in AppleShare IP. Anyone capable of setting up file sharing on a Macintosh should be able to set up AppleShare IP without assistance.

We've barely mentioned the use of dedicated servers in this book. That was intentional. You don't need to buy an expensive dedicated server and software to operate a small network with PCs or Macs or both. Because of the expense and added complications associated with the use of dedicated servers, and because they're not necessary (although they are powerful), servers aren't covered in this book.

You may think you don't need to bother with network security inside your own home or office. But are you sure you want your kids' friends or all your coworkers to have access to every shared file on your computer? There are various strategies for securing your shared files, some of which we've mentioned earlier in the chapter. Let's review.

- First, you don't need to turn on file sharing on every computer to move files back and forth. If you enable file sharing on your kids' computers (for instance) but not on your home office computer, you'll be able to move files back and forth from your computer, but the kids won't be able to accidentally open your investment portfolio (or other personal data) from their computers.
- You've seen that adding passwords on Windows machines and Macs isn't all that difficult. But you don't have to password-protect everything. If you want to share some files without a password, consider creating a special public folder and putting these files in it.
- You can choose to share individual folders rather than your entire hard disk.
- On a Mac, consider turning off the Owner's permission to see all disks that aren't shared items. This leaves fewer files at risk in the event someone learns your owner name and password. You set this permission by opening the owner icon in the Users & Groups control panel.
- If you have a cable modem connection to the Internet, you are at particular risk from your neighbors who also have cable modem connections. That's because all the cable modems in a neighborhood of several hundred cable TV customers are interconnected in one big local network. If your home network connects to the Internet via cable modem, be sure that all your computers with file sharing turned on require passwords, and make sure these passwords remain secure.



Security Reminders

Summary

File sharing eliminates the need to tote files from one computer to another on a floppy or Zip disk. Both Windows and the Mac OS let you designate folders and disks as shared items whose contents other people can access from other computers on your network. You then set access controls that determine who can read and write files on your computer. Windows and Mac OS have different procedures for setting up file sharing.

- ▼ In **Windows**, you must first **install the service** "File and printer sharing for Microsoft Networks" in the Network dialog box.
- ▼ You **control access in Windows** by assigning passwords to each shared folder or disk. Each folder and disk can have one password for Read-Only access and another password for Full access.
- ▼ You use the **Net Watcher utility** to see who's connected to a PC's shared folders and disks.
- ▼ On a **Macintosh**, the **file sharing software** is pre-installed, but you can turn it on or off.
- ▼ You control access on a **Mac OS** by setting **access privileges** for each shared folder and disk. You can set access privileges separately for three categories of network user: Owner, User/Group, and Everyone. You can set each category's privileges to Read & Write, Read only, Write only, or None.
- ▼ If you don't want to allow unidentified Mac users to connect without passwords, you can turn off all **Guest access**.
- ▼ To identify Mac network users and groups of users, you add them to the **Users & Groups** control panel.
- ▼ In Mac OS 8 and later, the **Activity Monitor** tab of the File Sharing control panel lists connected users and shared items. The same information appears in the **File Sharing Monitor** control panel in Mac OS 7.6.1 and earlier.
- ▼ Folders can **inherit access privileges** from the folder that encloses them, or each enclosed folder can have its own access privilege settings.

- ▼ **To share Mac files with Windows PC users**, you can use the Personal Web Sharing feature of Mac OS 8 and later. **To enable a PC to use Macintosh file sharing**, install Miramar Systems' PC MACLAN on the PC. To give a Mac the ability to use Windows file sharing, install Thursby Software's DAVE on that Mac.
- ▼ For the ultimate in cross-platform file sharing, hang a dedicated file server on your network. Windows NT Server is the most popular server software for PCs, as is AppleShare IP for Macs. Both serve files to Windows and Macintosh machines.
- ▼ File sharing makes your computers more **vulnerable to invasion**, especially if your network shares a cable modem connection to the Internet. You can take several steps to secure your shared files.

Want to use a network in your home or office but don't know where to start?

Learn just what you need to know to

- Set up and run a small home or office network
- Share an Internet access line
- Share printers and other equipment
- Share and access documents and files on other computers
- Play games and share calendars
- Back up and synchronize files
- Troubleshoot and manage your network
- And much, much more!

If you're interested in setting up a small computer network in your home or office, *The Little Network Book* will answer all your questions in a clear, no-nonsense style. Authors Lon Poole and John Rizzo explain networking hardware, software, setup, and protocols in simple, easy-to-understand terms to help you build useful, fun, and creative networks. Anyone who's new to networking or wants to improve a communication system at home or in the office will find *The Little Network Book* a valuable resource.




Peachpit Press

1249 Eighth Street • Berkeley, CA 94710

510-524-2178 • fax: 510-524-2221 • 800-283-9444

Find us on the World Wide Web at <http://www.peachpit.com>



Book Level:

- ☒ Beginning
- ☒ Intermediate
- ☐ Advanced

For Computers Using:

Windows 95/98
Macintosh OS 7.5.5 or higher

Computer Book Shelf Category:

Networking/Home or
Small Office

USA \$19.99
Canada \$29.95



ISBN 0-201-35378-4



ROKU EXH-1002

"The gospel of home LANs..."

David Wall, *Amazon.com Editorial Reviews*, www.amazon.com

This Wired Home

SECOND EDITION

The Microsoft® Guide to Home Networking

- Build a simple network for your home or home-based business
- Work smarter by sharing programs, printers, and an Internet account
- Play multiplayer interactive games—and take on challengers over the Net!

Alan Neibauer



ROKU EXH 1002

APPENDIX I

"A simple, understandable, and well-illustrated primer to the labyrinthine charms of home networking."

Paul Andrews, *The Seattle Times*

If you can plug in a PC, you can build your own home network!

Are there two or more computers in your house—but only one printer? Do your kids want to play games on the Internet at the same time you want to check your e-mail? Is your entire household competing for the same dial tone?

If you're running more than one PC under your roof, **THIS WIRED HOME** can show you how to build a simple network—and quickly multiply the computing power for your family or home office.

Just follow the easy step-by-step instructions for creating a secure and reliable network that can grow as your family or business grows. This how-to guide is written in plain, nontechnical language so you can put the information to work right away. You'll learn how to:

- Save time—save *money*—by sharing files, programs, printers, and other resources
- Match a networking solution to your needs—from direct PC-to-PC connections to DSL, wireless systems, and networks that run on your home electrical or phone lines
- Find the best free and low-cost connection software—including using the built-in capabilities in the Microsoft® Windows® Me, Windows 2000 Professional, Windows 98, and Windows 95 operating systems
- Hook up PCs and Macs on the same network
- Share a single phone line and Internet account
- Use your private network to send electronic sticky notes, maintain a central calendar, play multiplayer interactive games, and dial in from the road
- See what's ahead for home networking technologies and how to make wiring and setup decisions now that can evolve with the times

Like any home improvement project, all you need to build your own network is some guidance and the right tools. And with **THIS WIRED HOME**, you get the tips, tricks, and know-how to do it yourself!



About the Author:

Alan Neibauer has written several best-selling computer books, including *Running Microsoft Outlook® 2000* and *Small Business Solutions for Networking*. With a master's degree from Wharton, Alan has helped organizations of all sizes network their business information systems. He's also served as chairperson for an innovative computer MIS program at the university level.

Net

THIS WIRED HOME-002

NEIBAUER A3841 Netuk Gen G010
6309792 QP 1 10/16/01 MIGT
703-19H 000167963

2580

90000

U.S.A. \$29.99
U.K. £20.99
Canada \$43.99
[Recommended]

To learn more about
Microsoft Press® products, visit:
mspress.microsoft.com

Microsoft®
ROKU EXH. 1002

This Wired Home

SECOND EDITION

The Microsoft® Guide to Home Networking

Alan Neibauer

APPENDIX I

PUBLISHED BY

Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2000 by Alan R. Neibauer

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data

Neibauer, Alan R.

This Wired Home : The Microsoft Guide to Home Networking / Alan Neibauer.-- 2nd ed.

p. cm.

Includes index.

ISBN 0-7356-1158-0

1. Home computer networks--Amateurs' manuals. I. Title.

TK5105.75.N45 2000

004.6'8--dc21

00-057856

Printed and bound in the United States of America.

2 3 4 5 6 7 8 9 QWTQWT 5 4 3 2 1 0

Distributed in Canada by Penguin Books Canada Limited.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at mspress.microsoft.com. Send comments to mspinput@microsoft.com.

FrontPage, Microsoft, Microsoft Press, MSN, NetMeeting, Outlook, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

Unless otherwise noted, the example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

For Microsoft Press

Acquisitions Editor: Christey Bahn

Project Editor: Sally Stickney

For nSight

Project Manager: Susan H. McClung

Copy Editor: Chrisa Hotchkiss

Technical Editors: Don Lesser, Mannie White,
Doug Slaughter, Eric Brewer, Timothy Upton

You can also drag a document onto a printer icon that you've placed on the Windows desktop. To place a printer icon on the desktop, follow these steps in all versions of Windows:

1. On the Start menu, point to Settings, and then click Printers.
2. In the Printers window, right-click a printer and choose Create Shortcut from the shortcut menu.
3. When a message tells you that you can't place a shortcut in the Printers folder and asks whether you want to place the shortcut on the desktop instead, click Yes.

Connecting Printers Directly to the Network

Because a printer that's connected to a computer on the network works only when the computer is on, you might want to use an alternative: connecting the printer directly to the network. Connecting a printer directly to the network also frees up a computer's printer port so that you can hook up an external Zip drive, scanner, or other parallel device without a conflict.

In a twisted-pair network, you use twisted-pair cable to connect a printer to the hub. In a thin Ethernet network, you use coaxial cable to connect the printer to the network interface card (NIC) of the nearest networked device. Because the printer isn't connected to the printer port of a computer, anyone on the network can access it directly as long as the printer is turned on.

The disadvantage of connecting printers directly to the network is expense. Most printers are designed only for standard parallel connections. To connect them directly to the network, you'll need to purchase either a network-ready printer or a *print server*, a device that makes your printer network-ready.

Network-ready printers have a NIC built in. They cost more than standard printers and can be a little harder to find. The print server is equipped with an Ethernet connection on one side and a parallel, or possibly serial, connection on the other.

The least expensive print servers are called *pocket servers*. About the size of a pack of cigarettes, a pocket server plugs directly into a printer's parallel port. The twisted-pair cable from the network hub or the coaxial cable from another networked device plugs into the other end of the server.

Another type of print server connects to a printer with a cable. These external servers are usually more expensive than pocket servers, but they might include additional features. Some models, for example, have more than one parallel port, allowing them to connect several printers to the network at the same time.

Note

For some HP LaserJet printers, you can purchase an internal print server that fits inside the printer, much the way some NICs fit inside a computer.

When selecting a print server, make sure it matches your cable type—either twisted pair or coaxial. Some print servers, but not all, can accommodate both types.

The print server must also support the protocol you're using on your network. Some print servers support only IPX/SPX; others require either TCP/IP or NetBEUI.

Finally, while most printers have a standard-sized parallel port, called a *Centronics* port, some models, such as the LaserJet 1100, have a smaller mini-Centronics port. The standard-sized connection on a pocket print server won't fit a mini-Centronics port. If you're using such a printer, you'll need an adapter for the print server.

Note

To install an external print server, just connect the cable that came with the printer to the server's parallel connection. Connect the network cable to the server's network connection.

Setting Up a Pocket Print Server

Many different models of pocket print servers exist. Although they all operate in about the same way, their setup procedures vary. Most servers are sold with software that helps them connect to the network, but the process really depends on the type of protocol the server supports.

A TCP/IP server needs to be assigned an IP address. With a Windows peer-to-peer network, you'll probably have to assign the server a static IP address that isn't used by any computer on the network. Consequently, you might have to assign static IP addresses to every computer on the network as well, rather than have Windows assign them for you. Check the literature that came with your server for step-by-step directions for assigning it an IP address.

Most manufacturers provide programs to help you through the process. The Microplex Ethernet Pocket Print Server, for example, offers two programs for configuring the print server—IPAssign and Waldo. The IPAssign program, whose main dialog box is shown in Figure 11-5, accesses the print server through the Ethernet address and assigns it an IP address of your choosing.

**Figure 11-5.**

The IPAssign program for a Microplex print server assigns an IP address to the server.

The Waldo program is Java based, so you must have the Java runtime files installed on your computer. When you run Waldo, it searches for a Microplex print server on the network and displays its Ethernet address.

Device List						
Status	Ethernet Address	IP Address	Model	Serial	Version	Location
	00:80:72:03:21:ae		M205	08622	5.7	

You can then click the Assign button in the Waldo window to associate an IP address and subnet to the Ethernet address.



Once you assign an IP address to your server, you configure Windows to communicate with the printer. You first have to associate the server with a printer port. The default port most printers use is called LPT1, the parallel connector that the printer cable plugs into. When you configured your printer, as you learned in “Installing a Printer” earlier in

this chapter, you associated the printer with the port so Windows knows where to send the information to be printed—to the LPT1 port and then out to the printer.

When you connect a print server to the network, you need to create a port with which the IP address is linked. When you associate a printer to that port, Windows sends the information to be printed through the network and the Ethernet address of the print server.

How you associate a printer port to the print server depends on the print server itself. With Microplex servers, for example, the server appears as a device in Network Neighborhood or My Network Places and has four ports associated with it. When you configure the printer, you browse to the port you want to use in the same way you would browse to a workstation, as explained in the section “Accessing a Shared Printer,” earlier in this chapter.

Other manufacturers handle port assignments differently. The pocket print servers from Axis Communications, for instance, don’t appear in Network Neighborhood or My Network Places. Instead, you use the NetPilot program to associate the server with a port, and then you use a program called Axis Print System to add the printer to Windows.

Microplex and Axis certainly aren’t the only makers of pocket print servers. Table 11-1, later in this chapter, lists other print server makes and models.

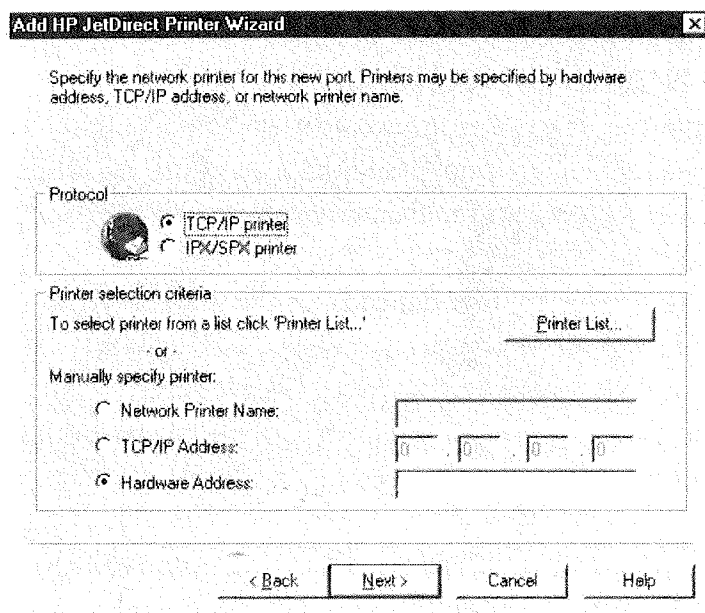
Setting Up an External Print Server

External print servers, an alternative to pocket print servers, connect to a printer by cable rather than plug directly into the printer itself. External servers work in the same way and are set up the same way as pocket print servers, although they’re more expensive than pocket print servers. Many models also come with two or more parallel connections that allow you to place multiple printers on the network so that you can use different printers for different documents.

Hewlett-Packard’s JetDirect print servers, for example, work with virtually any printer equipped with a parallel port—not just HP’s own brand. The line includes two models that have three parallel connections and a one-printer model, the 170X, that’s more suitable for home networks.

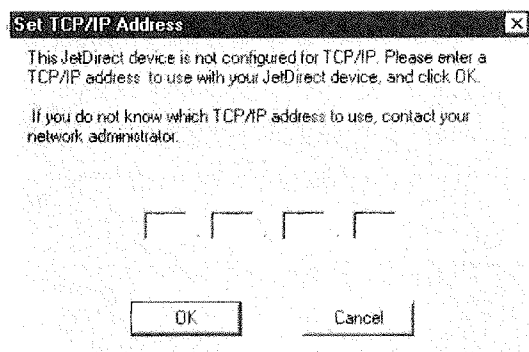
Setting up an HP print server is easy. After you connect the server both to the printer and to your network hub, you press a small button on the back of the server to print out a page of configuration information, including the electronic hardware address that is built into the device.

You then install the JetAdmin program supplied with the server and use the HP JetDirect Printer Wizard to configure the device. Figure 11-6 shows the wizard page in which you select a protocol and enter the unit’s hardware address.

**Figure II-6.**

The HP JetDirect Printer Wizard prompts you to select a protocol and enter the server's hardware address.

Using the address, JetAdmin locates the printer and displays a dialog box in which you can specify an IP address if you're using the TCP/IP protocol. After a few additional steps, JetAdmin starts the Add Printer Wizard in Windows, which opens a dialog box that prompts you to assign an IP address.



After the JetAdmin setup, you can send documents to the printer from your computer, and other network users can select the printer as their network printer and print documents even when your computer isn't on.

Many manufacturers of print servers exist, so you have plenty of choices. Table II-1 lists print server makes and models and each manufacturer's Web address.

Table 11-1. Print Server Manufacturers and Models

Manufacturer	Models	Web site
Axis Communications	Pocket, and one-port and two-port models, some with both parallel and serial ports	http://www.axis.com/
NETGEAR	One- and two-port models, some with built-in four-port hub	http://www.netgear.com/
Emulex	Pocket, and two-port and three-port models	http://www.emulex.com/
Extended Systems	Pocket, and one-port and two-port models, some with both parallel and serial ports	http://www.extendedsystems.com/
Hewlett-Packard JetDirect	One-port and three-port models, external and internal, and one model for sharing over home telephone lines	http://www.hp.com/
Intel NetPort Express	One-port and three-port models	http://www.intel.com/
Lantronix	Pocket and external print servers, up to six-port models (four parallel and two serial)	http://www.lantronix.com/
Linksys EtherFast	One-port and three-port models	http://www.linksys.com/
MicroPlex	Pocket, and a four-port model (two parallel and two serial)	http://www.microplex.com/


Sharing printers on a network can be a great time-saver and step-saver. You'll no longer need to carry a disk to another computer to print a document or carry a printer to another computer. With Windows, you don't have to purchase any additional software or hardware unless you want to connect your printer directly to the network.

Sharing files and printers isn't the only benefit of connecting computers on a network, however. You'll learn in the next chapter that you can use your network to create a family e-mail system for sending and receiving messages between family members.

ZIFF-DAVIS
A SOFTBANK
company

PC MAGAZINE

**PC Labs Tests 20 Digital Cameras
Plus: 18 Digital Imaging Packages**



- After Hours: 6 Encyclopedias on Disk
- 11 Servers to Let You Share Printers
- More Jim Seymour Win 98 Survival Tips

FIRST LOOKS:

- Exclusive: 250MB Iomega Zip Drive
- 19" Flat-Screen Monitors
- Diamond Rio PMP300
- Deneba Canvas 6
- Sun's Solaris 7.0
- Intel eMail Station

WWW.PCMAG.COM THE INDEPENDENT GUIDE TO PERSONAL COMPUTING VOL. 18 NO. 2 JANUARY 19, 1999

FREE SOFTWARE

ON THE WEB

THE BEST
DOWNLOAD SITES
FOR WORD PROCESSING,
WEB BROWSING,
E-MAIL, AND
MORE

PLUS
Tips and
Tricks for
Downloading

Digitized by Google

University of Michigan--Dearborn
Mardigian Library
PC magazine : the independent
guide to IBM-standard personal
computing
18:2
January 1999
Received from
UNIVERSITY OF MICHIGAN

ID***CAR-RT SORT**C006
00093 9#451472 1Q
LIBRARY NOV 30 99
NAZARENE CLG BALB
M AV #1467
MA 02170-2905



January 19, 1999
Volume 18
Number 2

UP FRONT

From the Editor-in-Chief4

Pipeline9

Letters21

Trends28

- More choices in wireless and phone-line home networking
- What's in an Oracle appliance?
- Tuning into desktop TV
- On-screen reading that's easy on the eyes
- Roam the planet
- Chips change their spots
- The Palm goes wireless

Inside PC Labs29

- USB speaker technology



First Looks41

- 250MB Iomega Zip Drive
- 19" flat-screen monitors
- Diamond Rio PMP300
- Daneba's Canvas 6
- PC radios from ICOM, Sony, and Ten-Tec
- O'Reilly Utilities for Windows 98 Annoyances
- Microtek ImageDeck
- Painter 5.5 Web Edition
- FoneCam
- L&H iTranslator Professional Service
- ALPS MD-5000
- LMSoft Presenter 3.0
- Systat 8.0

Networking First Looks75

- Sun's Solaris 7.0
- Intel InBusiness eMail Station

Second Looks80

- Norton AntiVirus 5.0
- Dragon NaturallySpeaking

COVER STORY


Free Software On the Web



System Utilities	100
Microsoft Office Add-Ons	102
Internet Tools	102
Communications Tools	104
Organizational Tools	104
Web Development Tools	114
Games	114
Tips on Web Downloading and Storage	101
Where To Get It: The Sites	103
Essentials	104

BY GAIL SHAFFER

Downloading software. Besides sending e-mail, it's still our favorite thing to do online—especially if it's free or almost free! We scoured the Web and found a bunch of great programs for you that do all sorts of things. And we surveyed the best download sites so you can save some time the next time you're looking for something.....100



Sharing Printers

BY STEVE RIGNEY

Still useful after all these years, print servers give you great flexibility in placing your shared printers right where you need them. They're inexpensive and easy to set up.....179



Axis Communications Inc.	185
Castelle Inc.	185
D-Link Systems Inc.	186
Emulex Corp.	186
Extended Systems Inc.	186
Hewlett-Packard Co.	186
Intel Corp.	186
Lexmark International Inc.	186
Linksys Inc.	187
Microplex Systems Ltd.	187
Oricom Technologies Inc.	187
Two Ways to Print	187
Editors' Choice	187
Single-Printer Sharing	187
Summary of Features	188
Performance Tests	188

FEATURES

CONSUMER IMAGE EDITING

Point-and-Shoot Software

BY SALLY WIENER GROTTA AND DANIEL GROTTA

You'll fall in love with your PC all over again once you start playing with digital images using one of these packages. Fix red-eye, adjust contrast and brightness—all with a click or two of the mouse—then move on to even more fun things, including making your own cards and calendars.....154



Adobe PhotoDeluxe Home Edition 3.0	156
Corel Print House Magic Deluxe 3.0	164
LivePix 2.0	167
MGI PhotoSuite II	167
Microsoft Picture It! 99	173
Ulead Photo Express 2.0	174
Editors' Choice	156
Suitability to Task	164
The Digital ShoeBox	166
Other Image Editors	174

PERIPHERALS

Digital Cameras

BY DANIEL GROTTA AND SALLY WIENER GROTTA

They've come a long way (in a short time), baby. The latest generation of digital cameras delivers near-35-mm quality photos in larger sizes. Models range from \$300 to \$1,200 and vary in features, so we've awarded Editors' Choices in three market segments. For all that digital cameras offer, they're worth the price.....116



Agfa ePhoto 780, ePhoto 1680	125
Casio QV-7000SX	125
Eastman Kodak DC210 Plus, DC260	125
Epson PhotoPC 700	128
Fujifilm MX-500, MX-700	128
HP PhotoSmart C30	131
Konica Q-M100V	131
Leica digilux	128
Nikon Coolpix 900s	148
Olympus D400 Zoom, D-620L	149
Ricoh RDC-4200, RDC-4300	149
Sanyo VPC-X300	150
Sony MVC-FD71, MVC-FD81	152
Toshiba PDR-M1	152
Editors' Choice	118
Performance Tests and Output Samples	133
Hitachi M2: Video, Too	148
Image Transfer: Easier than Ever	150
Summary of Features	151

Get more online at www.pcmag.com

Online Exclusive:

Photo Samples

How good do the photos taken by the digital cameras tested for this issue's story look? See for yourself at our Web site.

Plus: The Daily Download! A fresh new pick from the ZDNet Software Library every day.

Also Currently Online:

100 Desktop Fixes

Windows doesn't have to look that way. Customize it to your heart's content with this collection of hints and free utilities.

Plus:
▶ Test your PC for Y2K compliance live online! Are you ready?

Multimedia Extra: CD plus Web Site

Top Software

- ▶ 30 utilities ready to install
- ▶ Searchable back-issue database
- ▶ Web site link to more free utilities and the best downloads

Visit Extra at www.pcmagextra.com. To order, call 800-335-1195 in the U.S. or 303-665-8930 elsewhere.

See this issue's free utility in action at www.pcmag.com/download

COLUMNISTS

JAKE KIRCHNER	30
BILL MACHROME	85
JOHN C. DVORAK	87
INSIDE TRACK	89
JIM SEYMOUR	93
BILL HOWARD	99

SOLUTIONS

Tutor	194
Expansion buses provide vital connection to peripherals.	
User to User	197
Use Word to create presentations, automate Windows log-on, transpose Excel images easily, and more.	
Utilities	206
InCtrl4 makes keeping track of program installation easy.	
Internet User	212
You can get DHTML to work in both IE and Navigator.	
PC Tech	215
Microsoft Visual C++ 6.0 is both easier and more powerful.	

AFTER HOURS

A World View	247
Disk-based encyclopedias provide useful, up-to-date information.	



Shogo	252
A first-person shooting game that's in a league of its own.	



Also in This Issue:

PC Magazine Marketplace	226
Coming Up	254
Editorial Product Index	256
Advertisers' Index: Companies	257
Advertisers' Index: Products	260
Abort, Retry, Fail?	262

Get a DOWN LOAD of this.

W
W
W
.
P
C
M
A
G
.
C
O
M

250 FREE DOWNLOADS
(to be precise)



ZD
DAVIS



EDITOR-IN-CHIEF Michael J. Miller

EDITOR Jake Kirchner

SENIOR EXECUTIVE EDITOR Bill Howard

EXECUTIVE EDITORS Leon Erlanger, Peter McKie, Gail Shaffer, Don Willmet

DIRECTOR, PC LABS Steve Buehler

ART DIRECTOR Laura Baer

SENIOR NETWORKING EDITOR Frank J. Darfler, Jr.

SENIOR ONLINE EDITOR John Clyman **SENIOR EDITORS** Carol Venezia (First Looks), Sebastian Rupley (West Coast)

MANAGING EDITOR Paul B. Ross

SENIOR ASSOCIATE EDITORS Jamie M. Bailes (Hardware), Carol Levin (Trends), Sharon Terdeman (Technical Columns) **ASSOCIATE EDITORS** Eileen Bien (First Looks Networking), David Lidsky (Internet), John Morris (Software), Tom Ponzo (First Looks Online), Jennifer Triverio (After Hours), Anush Yeghazarian (PCs) **ASSISTANT MANAGING EDITOR** Kim Schueler **PRODUCTION EDITOR** Monica Sirignano **COPY CHIEF** Glen Boisseau Becker **STAFF EDITORS** Mary E. Bahr, Doug Belzer, Paul Dwyer, Matthew Graven, Nancy E. Hirsch, Josh Levy, Carol A. Mangis, Michael W. Muchmore **SENIOR WRITER** Cade Metz **STAFF WRITER** Angela Hickman **SENIOR COPY EDITORS** Jennifer Gollub, Joseph N. Levine **COPY EDITORS** Michael Feist, Jeremy A. Kaplan, Barbara McGeoch, Ann Ovodow **ASSISTANT COPY EDITOR** Sarah Pike **PRODUCTION MANAGER** Patricia Perkowski **SENIOR LAYOUT EDITOR** Lillian Gaffney **LAYOUT EDITOR** Michel Ologinski **PRODUCTION SYSTEM SUPPORT ANALYST** Nancy Goodman-Slayback **LIBRARIAN** Nancy Sirapyan **LIBRARY ASSISTANT** Dolores Williams **EDITORIAL RESEARCHERS** Adam Asch, Roderick A. Beltran, Richard Brown, Sharon Nash, Angela Tuka **DIRECTOR, IS TECHNOLOGY** Craig Ellison **NETWORK SUPPORT ANALYST** Melvin Acevedo **COMMUNICATIONS MANAGER** Anita Anthony **ASSISTANT TO THE EDITOR-IN-CHIEF** Christine Curti **ASSISTANT TO THE EDITOR** Rita Aghamian **ADMINISTRATIVE ASSISTANT** Christine Okang **SYSOP, PC MAGNET** Ken Hipple

CONTRIBUTING TECHNICAL EDITOR Neil J. Rubenking **CONTRIBUTING EDITORS** Bill Machrone (Vice President, Technology), Greg Alwang, Douglas Boling, Padraic Boyle, Bruce Brown, Sheryl Carter, John C. Dvork, Les Freed, Daniel Grotta, Sally Wiener Grotta, David Linthicum, Edward Mendelson, Jan Ozer, Charles Petzold, Stephen W. Plain, Alfred Poor, Jeff Prosser, John R. Quain, Neil Randall, Sal Ricciardi, Steve Rigney, Winn L. Rosch, Jim Seymour, Barry Simon, Luise Simone, Craig Stinson, M. David Stone

SENIOR ASSOCIATE ART DIRECTOR Lisa Kocurek **ASSOCIATE ART DIRECTORS** Eileen Hanley, Michael Scowden **ASSISTANT ART DIRECTOR** Lore Morgenstern **GRAPHICS DIRECTOR** David Foster **GRAPHIC ARTIST** Mark Tynner **ART PRODUCTION MANAGER** Talar Min **CONTRIBUTING PHOTOGRAPHER** Thom O'Connor **ASSISTANT TO THE ART DIRECTOR** Frieda T. Smallwood

TECHNICAL DIRECTORS Ben Z. Gottasman (Software), Nick Stam (Hardware), Jeffrey G. Witt (Networking and Communications) **DIRECTOR, OPERATIONS** John R. Delaney **SENIOR PROJECT LEADERS** Richard Fisco (PCs), Jay Munro (Internet, PC Tech) **PROJECT LEADERS** Laura Cox (First Looks), Russ Iwanchuk (Networking and Communications), Diane Jecker (Software), S. Jae Yang (Hardware), Kevin Young (Networking and Communications) **PROGRAMMERS** Richard V. Dragan, Win Swarr **PRODUCT TESTING MANAGER** Charles Rodriguez **TECHNICAL ANALYSTS** Oliver Kaven, Melanio Lopez, Jacqueline Paredes, Mark Valentine, Brad Walden, Martin Wong **SUPPORT TECHNICIANS** David Carela, William Pagan, Miriam Sampson, Sunita Sawh, Jeffrey Spada **INVENTORY CONTROL MANAGER** Tom Kennedy **INVENTORY CONTROL COORDINATORS** Richard P. Bifone, Bryan Hughes **ADMINISTRATIVE ASSISTANTS** Christina M. Evelyn, Leslie Sorich

PC MAGAZINE ONLINE

MANAGING EDITOR Tin Albano **SENIOR EDITOR** Josh Taylor **SENIOR TECHNICAL EDITOR** Thomas W. Giebel **TECHNICAL EDITOR** Webster T. Mudge **PROGRAMMER** Brian D. Buck **ASSOCIATE EDITOR** Edward Grossman **DESIGNER** Marcie Gandell **STAFF EDITORS** Jennifer L. Anderson, Troy Dreier, Richard Egan **PRODUCTION EDITOR** Matthew Slaybaugh

PUBLISHING DIRECTOR Nancy Newman

NATIONAL SALES DIRECTOR Vickie Pinsky **GROUP BUSINESS DIRECTOR** Bret Violette **ASSOCIATE BUSINESS MANAGER** Christine Holsten **RESEARCH DIRECTOR** Gordon Plutsky **MARKETING MANAGER** Dawn Gudelis

SENIOR ADVERTISING PRODUCTION MANAGER Ivis Fundichay **ADVERTISING PRODUCTION COORDINATORS** Milena Kotrch, Simone Oliver-Weekes **EDITORIAL PRODUCTION MANAGER** Pamela J. Berkowitz **MULTIMEDIA PRODUCTION MANAGER** Louise LaBerge

ADVERTISING OFFICE: One Park Ave., New York, NY 10016-5802; 800 33 MAG AD, 212-503-5100

THE INDEPENDENT GUIDE

PC Magazine is the Independent Guide to Personal Computing. Our mission is to test and review products and report fairly and objectively on the results. Our editors do not invest in firms whose products we review, nor do we accept travel tickets or other gifts of value from such firms. Except where noted, PC Magazine reviews are of currently available products. We review products without regard to advertising or business relationships with any vendor. Softbank, the majority holder of ZD Inc., has made a number of strategic investments in high-technology companies. A list of those companies is available online at www.pcmag.com/tag, and we will alert our readers to such investments whenever pertinent.

HOW TO CONTACT THE EDITORS

We welcome comments from readers. Send them to Internet address pcmag@zd.com or to PC Magazine, One Park Ave., New York, NY 10016-5802. Please include a daytime telephone number. PC Magazine's general number is 212-503-5255. The West Coast Operations number is 850-513-8000. We cannot look up stories from past issues, recommend products, or diagnose problems with your PC by phone. For an index of past issues and a list of upcoming stories, browse www.pcmag.com.

If you are dissatisfied with a product advertised in PC Magazine and cannot resolve the problem with the vendor, write (do not call) Ellen Askin, Advertising Department, at the above address. Please include copies of your correspondence with the vendor.

SUBSCRIPTION SERVICES

Internet: <http://subscribe.pcmag.com/service>
U.S. and Canada: telephone, 303-665-8930; fax, 303-604-7455
Elsewhere: telephone, 303-604-7445; fax, 303-664-0540
Mail: PC Magazine, P.O. Box 54853, Boulder, CO 80322-4093

Subscription rates. The one-year rate (22 issues) is \$49.97 in the U.S., \$85.97 elsewhere. Make checks payable to PC Magazine; U.S. currency only.

Back issues are \$8 each in the U.S., \$10 elsewhere (subject to availability). Prepayment is required. Make checks payable to PC Magazine; U.S. currency only. Mail your requests to Back Issues, ZD Inc., P.O. Box 53131, Boulder, CO 80322-3131.

Mailing lists. We sometimes make lists of our customers available to makers of goods and services that may interest you. If you do not wish to receive their mailings, please write to us. Include your mailing label with any correspondence; it contains information about your subscription that will facilitate processing. Please allow 6 to 8 weeks for your first issue to arrive or for any changes in your subscription to take place.

Additional information on advertised products can be requested online at www.pcmag.com/infolink.

PC MAGAZINE EXTRA

PC Magazine Extra, the interactive CD-ROM companion to PC Magazine, is available quarterly. To order (\$49.95 per year in the U.S., \$72 in Canada, \$99 elsewhere), please call 303-665-8930 in the U.S. or 303-604-7445 elsewhere. Or write to PC Magazine Extra, P.O. Box 54854, Boulder, CO 80322-9494.

PC MAGAZINE ONLINE

PC Magazine is on the World Wide Web (www.pcmag.com) and The Microsoft Network (Go to pcmagazine). We operate PC MagNet, an on-line service of ZDNet, hosted by CompuServe. (For details, see the Utilities column.) PC Magazine subscribers receive free access to our online archives.

PERMISSIONS, REPRINTS

Material in this publication may not be reproduced in any form without permission. If you want to quote from an article or use PC Magazine's logo in conjunction with an Editors' Choice designation, write Chantal Tucker or fax her at 212-503-5475; for information on reprints, please contact ZD Reprints at 800-825-4237.

The following are registered trademarks of ZD Inc.: DQMark, NetBench, PC, PC DIRECT, PC Labs, PC MAGAZINE, PC MAGAZINE AWARD FOR TECHNICAL EXCELLENCE, PC MagNet, ServerBench, WinBench, Winstone, and ZD.

The following are trademarks of ZD Inc.: Abort, Retry, Fail!, After Hours, BusinessCard, Corporate Developer, CPUmark, EasyComputing, Extending Your Apps, Features Plus, First Looks, First Looks Plus, Lab Notes, New & Improved, OFF THE STACK, PC Bench, PC Magazine At Home, PC Magazine CD, PC MAGAZINE EDITORS' CHOICE, PC Magazine Extra, PC Magazine Marketplace, PC Solutions, PC Tech, Pipeline, Power Programming, Quick Clips, Read Only, ScreenDemos, Tech Notes, Tutor, User-to-User, WinDraw, ZDNet, and ZiffNet.

(Other trademarks and trade names used throughout the publication are the property of their respective owners.)

Copyright (c) 1998 ZD Inc. All rights reserved. Reproduction in whole or in part without permission is prohibited.

Networking

BY
STEVE
RIGNEY

PRINT

SERVERS

Most departmental and enterprise printers come with their own network connections. If your printer doesn't come ready to share, however, you're not out of luck. One solution is to connect it to the departmental file server via a standard parallel cable. But since parallel cables become unreliable after some 25 feet, you may not be able to place the printer where it is most convenient to users. You can also connect your printers to client PCs in a peer-to-peer scenario, but in order to use those printers, you'll have to make sure the PCs are turned on. And the users of those client systems may suffer mysterious pauses while the PC routes print jobs.

The best way to connect network printers is through a print server, which is a small hardware device with a network connection, par-

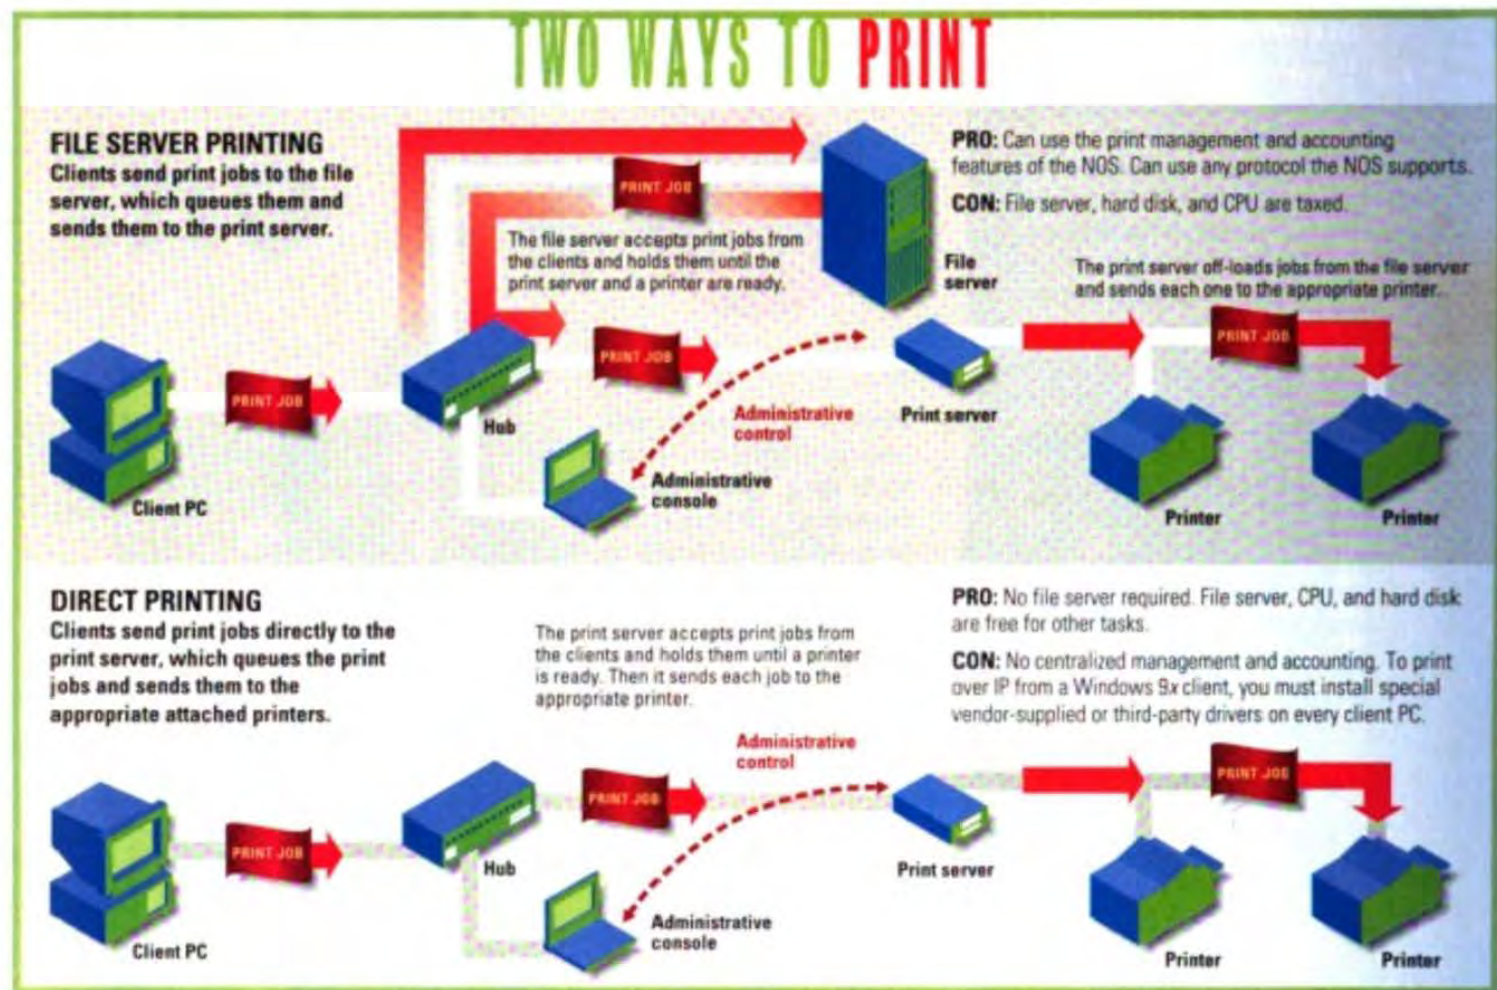
Today's print servers make sharing printers on the LAN easier and less expensive than ever before.

Axis Communications Inc.	185	Intel Corp.	189
Castelle Inc.	185	Lexmark International Inc.	189
D-Link Systems Inc.	186	Linksys Inc.	193
Emulex Corp.	188	Microplex Systems Ltd.	193
Extended Systems Inc.	188	Osicom Technologies Inc.	193
Hewlett-Packard Co.	189	Summary of Features	188

Digitized by Google

Original from
JANUARY 1999 PC MAGAZINE 179
UNIVERSITY OF MICHIGAN

NETWORKING Print Servers



allel and serial ports, and software that takes print jobs from clients or a server and routes them to the connected printers one at a time. Best of all, since print servers connect directly to the network, they let you place your printers just about anywhere you want.

PLAYERS AND PRICE

In this roundup, we look at 11 network print servers: the Axis PrintPoint 560/100 3P, Castle LANpress 3P/100, D-Link DP-300, Emulex NETQue PRO2, Extended Systems ExtendNet 100x, HP JetDirect 500X, Intel Netport-Express 10/100, Lexmark MarkNet Pro 3, Linksys EtherFast 3-Port 10/100, Microplex M202 Plus, and the Osicom NETPrint 1000 10/100.

The good news for small businesses and workgroups is that print servers have steadily come down in price since we

looked at them last year. For example, the dual-parallel-port D-Link DP-300 has a rock-bottom price of \$199.

The other good news is that we found these devices easier to set up than ever before. And most of them can be managed remotely either via Microsoft Windows software or over the Web with a standard browser.

WHAT TO BUY

All of the print servers we tested performed well. In fact, we found that the bottleneck in the printing process was not the print server but the printers to which they were connected. The differences among these devices lie in price, the number of printer ports they provide, whether they have Fast Ethernet or just Ethernet connections, and which operating systems and

protocols they support.

Another important difference lies in the type of remote management offered. All offer some kind of Windows-based management utility. Most now also come with Web servers and permit management over the Internet or company network via any standard Web browser. We especially liked the management packages from Extended Systems, HP, and Intel, which were thorough, well organized, and easy to use. HP's excellent JetAdmin

software lets you manage several types of SNMP devices and organize and map them logically by department, group, or capabilities.

Using HP's JetAdmin with HP printers, you also get monitoring capabilities that go beyond the typical basic errors and online/off-line notification to inform you graphically of paper jams, paper out, and the status of toner and other consumables. The same is true when you use the Lexmark print server management soft-

www.pcmag.com

To check out this feature story and our online archive of networking coverage, visit our Web site.

Our Contributors: STEVE RIGNEY is a contributing editor of *PC Magazine*. JEFFREY G. WITT is the technical director for networking and communications, and RYAN SNEDEGAR is a technical analyst at PC Magazine Labs. KEVIN YOUNG was the project leader and executive editor. LEON ERLANGER was in charge of this story.

NETWORKING Print Servers

ware with Lexmark printers. Interesting to note, the Axis and Extended Systems print servers also work with the HP Web JetAdmin utility.

The Axis PrintPoint 560/100 3P can work as a Novell Distributed Print System (NDPS) device, which lets NetWare clients search for printers on the LAN by location and capabilities (color, paper size, and so forth), as well as automatically install appropriate printer drivers.

All of the servers come with at least two parallel ports, which should be enough for most installations. If you want to connect a third or fourth printer in the same location, the Castelle, Extended Systems, HP, and Linksys products provide a third port. Most also provide serial ports, but these are no longer widely used.

Many of the units also provide 10/100 connections for both Ethernet and Fast Ethernet networks. Only the Emulex and Microplex products were limited to 10-Mbps Ethernet.

And most of the print servers will work over just about any network protocol you have, including AppleTalk, IPX, NetBEUI, and TCP/IP. But make sure you can use the network protocol you want with your NOS. For example, the D-Link DP-300 can print only over NetBEUI under Windows 95 and over IPX under NetWare. The LinkSys product works only over IPX with NetWare and doesn't work with any protocol besides IPX and TCP/IP. This may require you to reconfigure all your Windows 95 clients if they're presently configured for TCP/IP.

PRINTING WITH A NET

A new standard in the works from the Printer Working Group of the Internet Engineering Task Force will let you print to any compliant printer on the Internet, if you have rights to it. Called the Internet Printing Protocol, it will let printers on the Internet serve a similar function that fax machines serve today. Final approval is scheduled for mid-December 1998. Only the HP JetDirect 500X supports the IPP standard today, but HP did not have a client ready for review. IBM currently

EDITORS' CHOICE

• HP JetDirect 500X



All of the print servers we tested performed well and were reasonably easy to set up and configure, but the HP JetDirect 500X represents the best combination

of functionality, management, and price.

For a list price of only \$399, you get a print server with three parallel ports, as well as support for the emerging Internet Printing Protocol and HP's networked all-in-one devices that let you scan directly to print. You also get HP's excellent JetAdmin software which goes a step beyond most print server management packages. JetAdmin lets you manage many different types of SNMP-compliant printers and other HP devices and also organize and map

them graphically by department, location, or capabilities. Combine the JetDirect with some market-leading HP printers and you get detailed information and alerts on consumables such as paper and toner.

We also liked the Intel NetPort-Express 10/100, which also has excellent setup and management software, as well as the Axis PrintPoint 560/100 3P and the Extended Systems Extend-Net 100x, which were easy to use, and can be configured with HP's JetAdmin software in addition to their own Web-based management packages.

offers an IPP prototype client driver as a free download from its Web site (www.printers.ibm.com/ipp/ipp.html). Expect to see printer URLs on business cards in the near future.

Axis PrintPoint 560/100 3P

\$499 list. Woburn, MA; 800-444-2947, 781-938-1188; www.axis.com; 900 at www.pcmag.com/infolink.

● Last year, we liked the PrintPoint 560/100 3P's small footprint, good per-

formance, and easy setup. This year, Axis has added a TCP/IP driver for Windows 9x-based clients and the ability to serve as a Novell Distributed Print System device.

If you already have HP print servers on your network, you can use HP's popular JetAdmin software, which HP packages with its own print servers, to configure and monitor the Axis PrintPoint 560/100 3P. Axis also bundles its excellent

Windows-based NetPilot utility, which automatically looks for Axis print servers and guides you through setup. You can then use any browser to connect to the built-in HTTP server and manage the device.—Steve Rigney

Castelle LANpress 3P/100

\$429 list. Santa Clara, CA; 800-289-7555, 408-496-0474; www.castelle.com; 901 at www.pcmag.com/infolink.

● The most exciting feature of the Castelle LANpress 3P/100 is its Internet Printing package, which lets you produce output by sending e-mail



The Axis PrintPoint 560/100 3P has one of the better Web-based management packages we tested.

Digitized by Google

Original from
JANUARY 19, 1999 PC MAGAZINE 105
UNIVERSITY OF MICHIGAN

NETWORKING Print Servers

SINGLE PRINTER SETUP

The products in our main roundup offer ports for attaching multiple printers to a network. If you have only one printer to attach, you can go with a less expensive single-port server, but otherwise, the buying criteria are the same. You want a product that accommodates your operating systems and network protocols and that is simple to set up and manage. And if you're tight on space, you'll want one with a small footprint. We looked at four of these diminutive devices from Extended Systems, I-Data, Intel, and Linksys. All provide a single parallel port and work in Windows NT and NetWare networking environments. Any of these products will do the job. Your choice comes down to features and price.

EXTENDED SYSTEMS POCKETPRO

Though not much larger than the end of a printer cable, the Extended Systems PocketPro (\$190 street) packs plenty of features, including very good Web-based management. The parallel port of the product can attach directly to the back of your printer, and there's an RJ-45 connector for twisted pair wiring to your 10-Mbps network (but not 100 Mbps). The ExtendView software makes installation and management amazingly simple; printer-condition alerts were the easiest to set up of any product here.

(Extended Systems Inc.; 800-235-7576, 208-322-7576; www.extendsystems.com; 904 at www.pcmag.com/infolink.)

I-DATA EASYCOM ETH 100

Slightly smaller than a cigarette pack, the I-Data EasyCom Eth 100 (\$300 street) takes Web management to the hilt, shunning any additional software. After you plug the unit into the printer, it prints out installation instructions. As with the other



products reviewed here, you can assign IP addresses manually or via DHCP. You can set up e-mail alerts, but the instructions for doing so are sparse. (I-Data Inc.; 516-243-6600; www.i-data.com; 911 at www.pcmag.com/infolink.)

INTEL NETPORT EXPRESS 10/100

We found the installation software of the Intel NetPortExpress (\$260 street) to be

the best of the lot. The Netport, about the size of a paperback book, offers 10- and 100-Mbps connectivity, as well as the obligatory Web server for anywhere management. Besides the RJ-45 and parallel connections, the unit has a test button for troubleshooting and DIP switches if you wish to force the auto-sensing networking to either speed. For additional management, the Netport Manager software offers a drop-down menu that lets you configure alerts to work in conjunction with Intel's LANDesk Management Suite. (Intel Corp.; 800-538-3373, 503-264-7354; www.intel.com; 906 at www.pcmag.com/infolink.)

LINKSYS POCKETPRINT SERVER PPS1

The least expensive product we tested, the Linksys PocketPrint Server PPS1 (\$170 street), fits in your palm and accommodates 10-Mbps networks via a RJ-45 connector, plus a BNC connector for older, thin net coaxial cabling systems. The software is Windows only; there's no Web management. A Wizard-style setup covers the monitoring bases. One minor nit: to configure the unit for peer-to-peer networking, you'll need to install proprietary software on each client. (Linksys Inc.; 800-546-5797, 949-261-1288; www.linksys.com; 908 at www.pcmag.com/infolink.)—Jeffrey Witt



Castelle's LANpress 3P/100 lets you print over the Internet via e-mail.

over the Internet to the print server, thereby providing a capability similar to a fax machine. But though you can use this feature to print over the Internet, the LANpress is not compliant with the Internet Printing Protocol, as is the HP product. Otherwise, the LANpress is a competent three-port Fast Ethernet print server that works with multiple protocols and performs well.

We liked the LANpress's bundled MPAdmin configuration utility but not that it works only over IPX. The LANpress was also the only other print server besides the D-Link DP-300 that does not include Web-based management capability.—SR

D-Link DP-300

\$199 list. Irvine, CA; 800-326-1688, 949-455-1688; www.dlink.com; 902 at www.pcmag.com/infolink.

● The D-Link DP-300 is the least expensive multiport print server we tested. It comes with three ports and can work with most major protocols and OSs. Unfortunately, you don't get the same protocol flexibility as with most of the products we tested. Windows 9x clients can print only using NetBEUI, and NetWare clients are limited to IPX. Only Windows NT clients can print over TCP/IP. Most of the other products in this roundup provide print redirectors that let the client OS print

using almost any protocol the print server supports.

NETWORKING
Print Servers

SUMMARY OF FEATURES								
Print Servers								
	Axis PrintPoint 568/100 3P	Castelle LANpress 3P/100	D-Link DP-300	Emulex NETQue PRO2	Extended Systems ExtendNet 100x	PC MAGAZINE EDITOR'S CHOICE HP JetDirect 560X	Intel NetPortExpress 10/100	Lexmark MarkNet Pro 3
List price	\$499	\$429	\$199	\$449	\$479	Ethernet 10/100TX version, \$399; Token-Ring version, \$619	\$399	\$409
Street price	\$300	\$340	\$180	\$280	\$320	Ethernet 10/100TX version, \$320; Token-Ring version, \$500	\$330	\$330
Processor type	Axis 32-bit RISC controller	Intel 186	AMD 80186	Intel 186	Intel 386EX	Custom HP ASIC	Intel 486	Hitachi SH2 27MHz
Total memory	2MB	768K	640K	1MB	2MB	4MB	3MB	2MB
Flash memory/Buffer memory	512K / 256K	512K / 256K	512K / 128K	1MB / 64K	1MB / 1MB	2MB / 2MB	1 MB / 2MB	1MB / 256K
Parallel/serial/other ports	2/1/1	3/0/0	2/1/0	2/1/0	2/1/1	3/0/1	2/1/0	2/1/0
Dimensions (HWD, in inches)	1.0 x 2.2 x 4.7	1.0 x 9.1 x 5.2	1.0 x 7.4 x 4.5	2.3 x 9.5 x 5.4	2.0 x 8.0 x 6.5	1.2 x 11.0 x 5.0	1.3 x 6.6 x 4.4	1.5 x 10.5 x 5.4
10/100 auto-sensing Ethernet connection	■	■	■	■	■	■	■	■
BootP / DHCP / DNS	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■
WINS	■	■	■	■	■	■	■	■
LPD support	■	■	■	■	■	■	■	■
Network Operating Systems								
Windows NT, 98, 95, and 3.x	■	■	■	■	■	■	■	■
NetWare/Unix/Macintosh	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■
Network Communications Protocols								
TCP/IP / IPX / NetBEUI	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■	■ ■ ■
AppleTalk	■	■	■	■	■	■	■	■
Management Software								
Windows NT, 98, 95, and 3.x	■	■	■	■	■	■	■	■
DOS/Unix	■ ■	■ ■	■ ■	■ ■	■ ■	■ ■	■ ■	■ ■
Embedded Web server	■	■	■	■	■	■	■	■
E-mail notification of errors	■	■	■	■	■	■	■	■
Allows remote resetting of server	■	■	■	■	■	■	■	■
Includes an SNMP agent	■	■	■	■	■	■	■	■
Can be configured through a browser	■	■	■	■	■	■	■	■
Service and Support								
Warranty	3 years	2 years	Lifetime	2 years	5 years	3 years	3 years	3 years
Technical-support hours (eastern time)	9:00-5:15 M-F	9:00-8:00 M-F	9:00-9:00 M-F	24 hours, 7 days	7:00-6:00 M-F	24 hours, 7 days	10:00-8:00 M-F	9:00-9:00 M-F, noon-6:00 Sat.
N/A—Not applicable. The product does not have this feature.								

er can recognize. We would also like to see D-Link incorporate Web-based management into the DP-300. Still, we found this product easy to install and configure, and its performance was perfectly acceptable, though not top of the line.—SR

Emulex NETQue PRO2

\$449 list. Costa Mesa, CA; 800-368-5391, 714-662-5600; www.emulex.com; 903 at www.pcmag.com/infocenter.

● The Emulex NETQue PRO2 performs reasonably well, but it's one of only two products (Microplex is the other) that can print only over a 10-Mbps connection. The box also contains a BNC connector in case you happen to need the older interface.

There's no quick-start guide; all documentation is on the bundled CD or Emulex's Web site. Luckily for us, the unit was fairly easy to install. The NETQue PRO2 includes a built-in Web server for management and configuration, but we had to download a new version of the firmware to use this feature. And documentation for the Web server interface is online only; the company plans to add this information to the CD.—SR

Extended Systems ExtendNet 100x

\$479 list. Boise, ID; 800-235-7576, 208-322-7575; www.extendsystems.com; 904 at www.pcmag.com/infocenter.

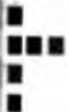






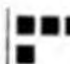




● The Extended Systems ExtendNet 100x is one of the most flexible and easily

installed units we reviewed. The small device has ports for four external printers, along with an auto-sensing 10/100-Mbps Ethernet connector. Note that you can use only three of the parallel ports out of the box. The fourth can work as a serial or parallel port but only with a special cable you buy from the vendor.

We liked the ExtendNet 100x's many configuration options. You can use the bundled ExtendWeb HTML utility or the Windows-based ExtendView for NetWare and TCP/IP. We were even able to use HP's popular JetAdmin and Web JetAdmin to set up print server options. The ExtendNet 100x was also the top performer under both NetWare and Windows NT.—SR

NETWORKING

Print Servers

Linksys EtherFast 3-Port 10/100	Microplex M202 Plus	Osicom NETPrint 1000 10/100
\$299	\$795	\$359
\$190	N/A	\$335
AMD 156	Motorola 68340	ARM Risc Processor
896K	3MB	2MB
512K / 128K	1MB / 16K	1 MB / N/A
3/0/1	2/2/0	2/1/0
10x90x55	13x8.6x5.4	20x7.0x4.0
		
		
		
		
1 year 10:00-8:00 M-F, noon-4:00 Sat	5 years 10:30-7:30 M-F	6 years 8:30-7:00 M-F

HP JetDirect 500X

\$399 list. Palo Alto, CA; 800-333-1917, 208-323-2551;
www.hp.com; 905 at www.pcmag.com/infolink.



PC
MAGAZINE
EDITORS'
CHOICE

Our Editors' Choice, the HP JetDirect 500X, comes with Hewlett-Packard's excellent Windows-based JetAdmin and Web-based Web JetAdmin management utilities. JetAdmin and Web JetAdmin automatically detected all of our print servers. Web JetAdmin can be used to monitor any HP or non-HP print server with its own Web server, and both utilities let you arrange and graphically map printers and other SNMP devices in logical groups by department, location, or

The screenshot shows the webJetAdmin interface. The top navigation bar includes icons for Home, Reports, Users, Groups, Settings, and others. The left sidebar contains a 'New Group' dropdown and a 'Check System Update' button. The main content area displays a table of Marketing items with columns for Service Address, Device Name, Status, and Description. The table lists several devices with their addresses, names, and current status (e.g., 'Try 1 Hour', 'Device Communication Error', 'Online').

HP's Web JetAdmin lets you manage several SNMP devices from a single console.

capabilities. You also get some extra monitoring features when you use JetAdmin with certain HP printers, including paper and toner status.

We were impressed with the product's ability to work with HP's networked all-in-one devices, which let you scan documents directly to the print server across the network. You can also use JetAdmin to configure print servers from Axis and Extended Systems and to manage HP CD-ROM servers.

The JetDirect is also the only print server we reviewed that claimed to comply with the Internet Printing Protocol draft. HP hadn't released its IPP client software in time for testing, however. —SR

Intel NetportExpress 10/100

\$399 list. Hillsboro, OR; 800-538-3373, 503-264-7354; www.intel.com; 906 at www.pcmag.com/infolink.

● The Intel NetPortExpress 10/100 is inexpensive, but it offers only two parallel ports and one serial port versus three parallel ports for the similarly priced HP and Linksys products. We especially liked the NetPortExpress's compact size (1.3 by 6.6 by 4.4 inches) and thorough management options. The package comes with excellent bundled Windows software or an SNMP management package and has

all you need for configuring and monitoring the unit using a browser.

Similar to HP's JetAdmin, Intel's management software lets you monitor and group multiple print servers. Its management capabilities, however, do not extend to other Intel and non-Intel devices.

Intel claims that with a 486 processor and 2MB buffer, the Net-portExpress is the fastest print server on the market. It wasn't the top per-

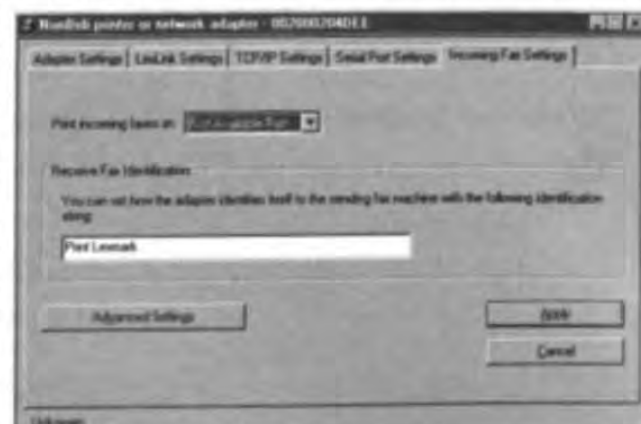
former on our tests, which found printers to be the performance bottlenecks in a real-world environment, but it performed very well.—SR

Lexmark MarkNet Pro 3

\$409 list. Lexington, KY; 800-539-6275, 606-232-2000; www.lexmark.com; 907 at www.pcmag.com/infolink.

● This is the only product in this roundup that lets you attach a fax modem to its serial port and receive incoming faxes (you can't send a fax). Used together with Lexmark printers and the included management utilities, you can also receive detailed alerts via graphical representations of the connected printers, making the MarkNet a must in networked Lexmark printer environments. Otherwise, the Lexmark is typical of the products reviewed here, with two parallel ports and good protocol support.

We liked the detailed, multilingual documentation, but we found the bundled



The Lexmark MarkNet Pro 3 lets you receive faxes via a fax modem attached to the serial port.

Original from
JANUARY 19, 1990 PC MAGAZINE 180

NETWORKING Print Servers

BENCHMARK TESTS

Print Servers



All the products we tested performed well. The printer, not the print server, typically is the performance bottleneck.

We used the fastest printers we could find—a 34-ppm Lexmark Optra S 3455 with 8MB of RAM and a 17-ppm HP LaserJet 4000 with 24MB of RAM—in order to make the print server work as hard as possible. Still, we found only a 20 percent variation between the fastest and slowest products. Buffer size, processor speed, and the Ethernet port speed had no significant effect on performance. For example, the Intel NetportExpress 10/100 featured an Intel 486 chip, and the Extended Systems ExtendNet 100x had a 386; yet they had similar output times.

Pooling the printers together in a single queue let Windows NT and NetWare send most of the print jobs to the Lexmark Optra, the faster printer. This resulted in faster printer performance than when we divided the print jobs into two queues, which used the slower printer, the HP LaserJet 4000, more often. We were forced to print results from two queues in Windows NT, however, because single-queued results varied significantly from one test to another. The Microplex M202 Plus in particular had problems handling Windows NT, at first taking over an hour to print using two print queues. Microplex suggested we abandon its driver and try the generic Windows LPR ports instead, which remedied the problem.—*Analysis written by Martin Wong*

How We Tested

We tested the print servers using a variety of test files, including 6-page Excel, 5-page Word, 35-page Word, 5-page PDF, and 10MB TIFF files. We tested under Novell's IntranetWare 4.11 with Support Pack 5B, and Windows NT with Service Pack 3. We installed the NOSs on identical Compaq 1600R servers with dual 266-MHz Pentium II CPUs and 512MB of RAM. Our clients were eight Dell OptiPlex GXas with Pentium II/233-MHz processors and

PRINT SERVER THROUGHPUT

Low scores are best. Minutes: seconds.
Bold type denotes first place within each category.
PC denotes Editors' Choice.

NETWARE 4.11	
Axis PrintPoint 560/100 3P	5:05
Castelle LANpress 3P/100	5:34
D-Link DP-300	4:51
Emulex NETQue PRO2	5:53
Extended Systems ExtendNet 100x	4:42
HP JetDirect 500X	5:00
Intel NetportExpress 10/100	5:05
Lexmark MarkNet Pro 3	4:58
Linksys EtherFast 3-Port 10/100	5:35
Microplex M202 Plus	5:14
Osicom NETPrint 1000 10/100	4:50

WINDOWS NT SERVER 4.0	
Axis PrintPoint 560/100 3P	6:47
Castelle LANpress 3P/100	6:44
D-Link DP-300	6:36
Emulex NETQue PRO2	6:40
Extended Systems ExtendNet 100x	6:19
HP JetDirect 500X	6:37
Intel NetportExpress 10/100	6:43
Lexmark MarkNet Pro 3	6:35
Linksys EtherFast 3-Port 10/100	6:48
Microplex M202 Plus	6:41
Osicom NETPrint 1000 10/100	6:44

32MB of RAM, eight Dell Dimension XPS Pro200ns with Pentium Pro/200 processors and 32MB of RAM, and eight Dell Dimension XPS P166s with Pentium/166 processors and 32MB of RAM, each running Windows 95. The PCs were connected by SynOptics 28115 Fast Ethernet switches. A connection through a 3Com Superstack II Hub 100TX let us capture network data.

For NetWare, we printed using NDS in Pconsole mode. We used TCP/IP native drivers if available; otherwise, we used LPR. The printers on the test-bed were a 34-ppm Lexmark Optra S 3455 with 8MB of RAM and a 17-ppm HP LaserJet 4000 with 24MB of RAM.

When we printed to two queues, we used each printer's proprietary driver, but when we pooled the printers into one queue, we used the standard HP PCL 5 driver. In Windows NT, we were able to pause the queue and load up all of the print jobs. We measured the time it took from removing the print-queue pause to finishing the final output page on the last printer finished. For NetWare, we started timing when the clients clicked the Print button and ended with the completion of the last page. For all of the tests, we used a protocol analyzer to capture the network traffic between the print server and the NOS client.

No Monthly
Payments

freewwwweb

ONE
TIME
FEE
\$99.95

Internet Access

E-Mail/NewsGroups

Nationwide Access

Freewwwweb is the

Internet Access

provider who doesn't

make you pay

month after month!

Call or Download Now

1.888.970.FREE

www.freewwwweb.net

Original from

Digitized by Google

UNIVERSITY OF MICHIGAN

NETWORKING

Print Servers

Windows- and Web-based management software a little less easy to use than HP's JetAdmin.—SR

Linksys EtherFast 3-Port 10/100

\$299 list. Irvine, CA; 800-546-5797, 949-261-1288; www.linksys.com; 908 at www.pcmag.com/infolink.

● At \$299 list (\$190 street), the Linksys EtherFast 3-Port 10/100 is the least expensive print server we tested that offers three parallel ports and a 10/100-Mbps connection. Unfortunately, the EtherFast works only over IPX and TCP/IP, and it lacks browser-based management. Unless you're looking to save a quick \$100, you'll find better flexibility in products from HP, Intel, or Extended Systems.

The bundled Bi-Admin program is not as detailed as those of other packages, but it let us monitor multiple devices and configure our EtherFast as both a NetWare and a Windows NT-based print server. Bi-Admin also automatically detects any EtherFast devices on your network without any prior configuration. Like the D-Link DP-300, the EtherFast works only

with IPX for NetWare clients, but unlike its competitor, the EtherFast lets you print from Windows 9x using TCP/IP.—SR

Microplex M202 Plus

\$795 list. Burnaby, B.C., Canada; 604-444-4232, 800-665-7798; www.microplex.com; 909 at www.pcmag.com/infolink.

● The Microplex M202 Plus is the only product that shipped with fiber-optic connectors, yet it's also one of only two products that can only connect to the network at 10-Mbps Ethernet. And if you opt for an AUI connector instead of fiber, the price is still a whopping \$575 (list)—far higher than any other unit in this roundup. Unfortunately, the price is not justified by features.

We ran into some performance problems when using the bundled Windows NT driver. Microplex recommended that we use the LPR printer driver with Windows NT, which brought performance in line with the other products we tested. The included configuration utility is easy to use but not on a par with that of the HP or Intel products.—Ryan Snedegar

Osicom NETPrint 1000 10/100

\$369 list. Waltham, MA; 800-243-2333, 781-647-1234; www.osicom.com; 910 at www.pcmag.com/infolink.

● Though not a standout, the Osicom NETPrint 1000 10/100 works as advertised. The setup utility doesn't automate NetWare configuration as much as most of the other products we tested, but setup is not overly complicated. The Osicom is the only product here with no buffer, so you need to print via a file server.

Osicom comes with browser-based management and includes an IPX-based browser for networks that don't use IP. At 100 Mps, the NETPrint unit we received could not use DHCP or BootP for an IP address assignment, but Osicom claims the problem will be fixed by the time you read this. The product comes with only a single LED to display status; we prefer the multiple LEDs you get with the Lexmark and Microplex products.

The NETPrint performed well, particularly with NetWare.—RS

The End of Networking As We Know It...

DUAL-SPEED HUBS 8 AND 16 PORTS
WITH BUILT-IN SWITCH PORTS

UNICOM has developed hubs that combine either an 8 or 16 port dual-speed hub with either a 1 or 2 port switch for half or full duplex operation.

The bottom line...
YOU SAVE BIG!

Now you don't have to buy an extra switch to expand your network. With UNICOM's dual-speed hubs, you can enjoy the benefits of an 8 or 16 port hub with built-in switching capabilities for 10Base-T and 100Base-TX networking. The 16-port hub has a built-in module bay for greater flexibility allowing the installation of an additional 10Base-T/100Base-TX or 100Base-FX switch module. These dual-speed hubs from UNICOM also allow you to extend the distance between hub connections by up to 100 meters (using the extra switch ports).

UNICOM is dedicated to bringing premium quality network hardware at down-to-earth prices.



Ask your dealer for this convenient new product. Or call

1-800-346-6668

Ask for Operator 106

UNICOM

City of Industry, CA
WWW.UNICOMLINK.COM

Digitized by Google

Circle 180 at www.pcmag.com/infolink

Original from
UNIVERSITY OF MICHIGAN

#1

Bestselling
Computer
Book SeriesOne of PC World's Top 10 Bestselling
Books — Updated for Windows 98!Over
1 Million
Copies Sold
Worldwide

PCs FOR DUMMIES®

6TH EDITION

A Reference for the Rest of Us!

by Dan Gookin

Author of *DOS For Dummies*®, 3rd Edition




**The Fun and Easy Way™
to Tame a Computer**

**Your First Aid Kit®
for Selecting Hardware
and Installing Software**

**Windows 98 and the
Internet — Explained
in Plain English**

ROKU EXH. 1002



by Dan Gookin



IDG Books Worldwide, Inc.
An International Data Group Company

Foster City, CA ♦ Chicago, IL ♦ Indianapolis, IN ♦ New York, NY

APPENDIX K

PCs For Dummies®, 6th Edition

Published by
IDG Books Worldwide, Inc.
An International Data Group Company
919 E. Hillsdale Blvd.
Suite 400
Foster City, CA 94404
www.idgbooks.com (IDG Books Worldwide Web site)
www.dummies.com (Dummies Press Web site)

Copyright © 1998 IDG Books Worldwide, Inc. All rights reserved. No part of this book, including interior design, cover design, and icons, may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording, or otherwise) without the prior written permission of the publisher.

Library of Congress Catalog Card No.: 98-87442

ISBN: 0-7645-0435-5

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

6B/SR/QZ/ZY/IN

Distributed in the United States by IDG Books Worldwide, Inc.

Distributed by Macmillan Canada for Canada; by Transworld Publishers Limited in the United Kingdom; by IDG Norge Books for Norway; by IDG Sweden Books for Sweden; by Woodslane Pty. Ltd. for Australia; by Woodslane (NZ) Ltd. for New Zealand; by Addison Wesley Longman Singapore Pte Ltd. for Singapore, Malaysia, Thailand, Indonesia and Korea; by Norma Comunicaciones S.A. for Colombia; by Intersoft for South Africa; by International Thomson Publishing for Germany, Austria and Switzerland; by Toppan Company Ltd. for Japan; by Distribuidora Cuspide for Argentina; by Livraria Cultura for Brazil; by Ediciencia S.A. for Ecuador; by Ediciones ZETA S.C.R. Ltda. for Peru; by WS Computer Publishing Corporation, Inc., for the Philippines; by Unalis Corporation for Taiwan; by Contemporanea de Ediciones for Venezuela; by Computer Book & Magazine Store for Puerto Rico; by Express Computer Distributors for the Caribbean and West Indies. Authorized Sales Agent: Anthony Rudkin Associates for the Middle East and North Africa.

For general information on IDG Books Worldwide's books in the U.S., please call our Consumer Customer Service department at 800-762-2974. For reseller information, including discounts and premium sales, please call our Reseller Customer Service department at 800-434-3422.

For information on where to purchase IDG Books Worldwide's books outside the U.S., please contact our International Sales department at 650-655-3200 or fax 650-655-3297.

For information on foreign language translations, please contact our Foreign & Subsidiary Rights department at 650-655-3021 or fax 650-655-3281.

For sales inquiries and special prices for bulk quantities, please contact our Sales department at 650-655-3200 or write to the address above.

For information on using IDG Books Worldwide's books in the classroom or for ordering examination copies, please contact our Educational Sales department at 800-434-2086 or fax 317-596-5499.

For press review copies, author interviews, or other publicity information, please contact our Public Relations department at 650-655-3000 or fax 650-655-3299.

For authorization to photocopy items for corporate, personal, or educational use, please contact Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, or fax 978-750-4470.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: AUTHOR AND PUBLISHER HAVE USED THEIR BEST EFFORTS IN PREPARING THIS BOOK. IDG BOOKS WORLDWIDE, INC., AND AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS BOOK AND SPECIFICALLY DISCLAIM ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. THERE ARE NO WARRANTIES WHICH EXTEND BEYOND THE DESCRIPTIONS CONTAINED IN THIS PARAGRAPH. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES OR WRITTEN SALES MATERIALS. THE ACCURACY AND COMPLETENESS OF THE INFORMATION PROVIDED HEREIN AND THE OPINIONS STATED HEREIN ARE NOT GUARANTEED OR WARRANTED TO PRODUCE ANY PARTICULAR RESULTS, AND THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY INDIVIDUAL. NEITHER IDG BOOKS WORLDWIDE, INC., NOR AUTHOR SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES. FULFILLMENT OF EACH COUPON IS THE RESPONSIBILITY OF THE OFFEROR.

Trademarks: All brand names and product names used in this book are trade names, service marks, trademarks, or registered trademarks of their respective owners. IDG Books Worldwide is not associated with any product or vendor mentioned in this book.



is a trademark under exclusive
license to IDG Books Worldwide, Inc.,
from International Data Group, Inc.

Chapter 1

Say Hello to Mr. Computer

In This Chapter

- ▶ Understanding what a computer is
- ▶ Understanding what a computer is not
- ▶ Knowing what a PC is
- ▶ Recognizing basic computer hardware
- ▶ Finding drive A
- ▶ Reading PC hieroglyphics
- ▶ Getting to know software
- ▶ Discovering a final PC fact

Who ever said that computers are easy to use? They can be fun. They can be aggravating. They can be enlightening. And they can most certainly be intimidating, which is where all the *dummy* stuff comes about. It's just too bad that computers don't pop out of the box, shake your hand, and give you a big hug. If that were true, this book wouldn't be necessary.

In a way, computers are like babies. They come packaged with great potential and, with the right care and handling, will achieve wondrous things and make you very proud. Like babies, computers require you to get to know them, to learn their moods and which buttons to push. It's a mutual relationship — and one that, fortunately, doesn't involve any drooling or diapers to change.

It's Just Another Electronic Gadget

A computer is that thing on your desk that looks like a TV set illegally parked by a typewriter. Call it whatever you like; it's basically a computer. But, because you may also have a computer on your wrist, in your car, or in the toaster, a more specific term is required: What you have on your desk is really a *PC*, a *personal computer*.



- ✓ Computers are essentially calculators with a lot more buttons and a larger display. They organize. They help you work with words and numbers. They can educate and entertain.
- ✓ Computers are not evil. They harbor no sinister intelligence. In fact, when you get to know them, they're rather dumb.
- ✓ Computers have the potential to be very friendly. Because you can read information on-screen, many computers give you a list of options, provide suggestions, or tell you what to do next. The microwave oven can't do that. Or maybe it can but refuses to.
- ✓ Computers don't flash 12:00 after a power outage.
- ✓ Perhaps the most important thing to remember about a computer is that *you* are in the driver's seat. You tell the computer exactly what to do, and it does it. The problem is that a computer obeys your instructions no matter what — even when you tell it to do something goofy. The art of dealing with a computer is a precise one.
- ✓ Please refrain from whacking your electronics.

What is a computer?

Computers defy description. Unlike other tools that have definite purposes, a computer can do a number of different things, solving an infinite number of problems for an infinite number of people. Just about anything that can be done with words, numbers, information, or communication can be done with a computer.

In a way, a computer is just another electronic gadget. Unlike the toaster and your car's carburetor, which are programmed to do only one thing, a personal computer can be *programmed* to do a number of interesting tasks. It's up to you to tell the computer what you want it to do.

- ✓ The computer is the chameleon of electronic devices. Your phone can be used only as a phone, your VCR only records and plays videos, and your microwave oven can only zap things (food, mostly). But a computer's potential is limitless.
- ✓ Computers get the job done by using *software*. The software tells the computer what to do.
- ✓ No, you never have to learn programming to use a computer. Someone else does the programming, and then you buy the program (the software) to get your work done.
- ✓ Your job, as computer operator, is to tell the software what to do, which then tells the computer what to do.

- ✓ Only on cheesy sci-fi shows does the computer ever tell *you* what to do.
- ✓ You can always *verbally* tell the computer what to do with itself. This happens millions of times a day, by programmers and nonprogrammers alike.
- ✓ Software is only half of the computer equation. The other side is *hardware*, which is covered in the following section.
- ✓ Computers can't clean up the house; they lack eyeballs and arms and legs. What you need if you want your house cleaned is a *robot*, which scientists haven't yet perfected for domestic use. When they do, buy *Robots For Dummies*.

What is not a computer?

Theatre. Real estate. Most livestock. Durable goods. Books and music. Fine art. The Hair Club for Men. Sushi. Inflatable stuff. Anything with a knob. Cat litter. False teeth. Everything in the Potpourri category on *Jeopardy!* Dolphin-safe tuna. Anything with grease on it. Anyone who's been on Larry King's show. Larry King himself. Larry King's dog.

What then is a PC?

PC means *personal computer*. It's the name IBM gave to its first personal computer, the IBM PC. That one computer was the Model T of all computers. Even though it was made in 1981, many of its design elements are included in today's models — much to the frustration of PC owners and manufacturers everywhere!



Some history you don't have to read

Not long ago, people referred to personal computers as *microcomputers*. This term came from *microprocessor*, the computer's main chip. The big "I want to control the world and foul up your phone bill" computers were called *mainframes*. Smaller, corporate- and college-size computers (which only fouled up paychecks or grades) were called *minicomputers*.

According to the geeks who ran the mainframes and minicomputers, *microcomputers* were hobbyists' playthings — toys. However, the features available on the personal computer — the microcomputer — that you can have on your desk today exceed many of the features of the early mainframes. So there.

- ✓ The term *clone*, and later *compatible*, was once used to describe any computer that used IBM PC-like hardware and could run PC software. Those terms are rarely used today because the standard PC has become more generic. In fact, computers are now sold based on which type of operating system they support; see the section “The operating system (Or “Who’s in charge here?”),” later in this chapter.
- ✓ The only personal computer that’s not a PC is the Macintosh. Its owners prefer to call it a *Mac* rather than a PC, even though the Macintosh is a personal computer.
- ✓ By the way, this book doesn’t cover the Macintosh. If that type of computer is more to your liking, rush out and buy *Macs For Dummies*, 6th Edition (published by IDG Books Worldwide, Inc.), by my pal, Computer Magician to the Stars, David Pogue.
- ✓ In addition to the Macintosh, several other variations on the personal computer exist. See Chapter 31 for my derogatory opinions.
- ✓ The original PC wasn’t created to start a dynasty, even though it did. Had its designers known how successful it would be, things may have been different (meaning *worse!*).

Hardware and Software

Computers have two parts: hardware and software.

Hardware is the physical part of a computer, anything you can touch. Hardware is nothing by itself but potential. It needs software to tell it what to do. In a way, hardware is like a car without a driver or a symphony orchestra without music (some orchestras are better that way, but I digress).

Software is the brains of a computer. It tells the hardware what to do and how to work. Without the software directing things, the hardware would just sit around and look formidable. You must have software to make a computer go. In fact, software determines your computer’s personality.

- ✓ Computer hardware isn’t anything you find in your local TrueValue store. With a computer, the hardware is the physical part — the stuff you can touch, feel in your hand, drop on the floor, lug through an airport, or toss out a window.
- ✓ Computer software is the brains of the operation — the instructions that tell the computer what to do, how to act, or when to lose your monthly report.



- ✓ Computer software is more important than computer hardware. The software tells the hardware what to do.
- ✓ Although computer software comes on disks (CDs or floppy disks), the *disks* aren't the software. Software is stored on disks just as music is stored on cassettes and CDs.
- ✓ Without the proper software, your computer is a seriously heavy paperweight.

Your Basic Hardware (A Nerd's-Eye View)

Figure 1-1 shows what a typical computer system looks like. I've flagged the most basic computer things you should identify and know about. They're just the basics. The rest of this book goes into the details.

Monitor: This TV set-like thing typically perches on top of the *console*. The glass part of the monitor is the *screen*, which is where the computer displays information or offers you insults or rude suggestions. Monitors are covered in detail in Chapter 13.

Console: The main computer box is called the *system unit* by geeky types. It contains your computer's guts on the inside, plus a lot of interesting gizmos on the outside. See the section "Stuff on the console," later in this chapter, for information about the greeblies pasted on the console.




Figure 1-1:
Basic
computer
things.

Keyboard: It's the thing you type on. La-de-da. Chapter 15 cusses and discusses the computer keyboard.

Mouse: It's not a fuzzy little rodent, but rather a computer mouse. These things are especially helpful in using all that graphical software out there. By the way, the mouse is pointing the wrong way in the figure. I did that so that you can see the two mouse buttons. See Chapter 14 for information about proper mouse orientation and button info.

Speakers: Most PCs can beep and squawk through their own speaker. To augment that noise, most computers now sold include a special sound card (that *multimedia* thing). Along with the sound card comes an additional set of stereo speakers, which most people plunk down on each side of the monitor. Pay more money, and you can even get a subwoofer to sit under the desk. Now *that* will scare the neighbors.

Printer: It's where you get the computer's output: the printed stuff, also called *hard copy*. Sashay off to Chapter 16 to increase your PC printer knowledge.

Connectors: Behind your computer are a bunch of holes, each of which is named Jack. Into these *jacks* you plug various peripherals and other devices the computer controls — such as the printer, as shown in Figure 1-1. Some parts of Chapter 11 are devoted to the various PC connectors on a computer's rump.

Lots of ugly cables: One thing they never show you — not in any computer manual and especially not in advertisements — is the ganglia of cables that live behind each and every computer. What a mess! These cables are required in order to plug things into the wall and into each other. No shampoo conditioner on earth can clean up those tangles.

- ✓ These parts of the computer are all important. Make sure that you know where the console, keyboard, disk drive, monitor, and printer are in your own system. If the printer isn't present, it's probably a network printer sitting in some other room.
- ✓ Part IV of this book covers computer hardware in detail.
- ✓ A computer really exists in two places. Most of the computer lives inside the console. Everything else, all the stuff connected to the console, is called *peripherals*. See Chapter 18 for more information about peripherals.

Stuff on the console (front)

The console is the most important part of your computer. It's the main thing, the Big Box. Every part of your computer system either lives inside the console or plugs into it. Figure 1-2 shows what a typical PC console may look like. I've flagged the more interesting places to visit, although they may appear in a different location than shown in the figure.

CD-ROM or DVD drive: This high-capacity disc looks exactly like a musical CD, although it contains computer information. Chapter 7 covers how you use and abuse CD-ROM and DVD drives and discs.

Future expansion: It's usually a blank spot on the front of your computer that enables you to add even more junk later. Such a space may already be taken on the computer, filled with such goodies as a tape backup unit, ZIP drive, another CD-ROM drive, another hard drive, or a mystery grab-bag assortment of other computer things many folks enthusiastically spend their hard-earned money on.




Figure 1-2:
Important
doodads on
the front of
the console.

Floppy drive: This slot eats floppy disks. Most PCs sold in the past few years come with only one floppy drive, dubbed drive A. Older PCs may have two floppy drives, and some very old models may even have the wider 5¹/₄-inch type. Part III of this book uses several chapters to discuss all the whatnots of disks and drives. It's a big spinnin' deal.

Air vents: Okay, this one isn't truly important, but most consoles sport some type of air vent on the front. Don't block the air vents! The thing has gotta breathe.

Buttons and lights: Most of the computer's buttons are on the keyboard. A few of the more important ones are on the console, and these buttons on fancier PCs are accompanied by many impressive tiny lights. These buttons and lights include the following:



On-off button: The PC's main power button, the one you use to turn the darn thing on. The on-off button is usually accompanied by a light, although computers make enough racket that you can usually hear when they're turned on.



Reset button: Allows you to restart the computer without going through the bother of turning it off and then on again. Chapter 4 explains why anyone in his right mind would want to do that.



Sleep button: A feature on some newer PCs and most laptops. Pressing this button causes your PC to go into a coma, suspending all activity without turning the computer off. Read all about this trick in Chapter 4.



Hard drive light: Flashes when the hard drive is working. Because the hard drive lives inside the console, this light is your reassurance that it's alive and happy and doing its job.

Older PCs may sport two additional features on the front side of the console. Typically, these doodads are for decoration only; their original function was important only to certain computers in the mid-1980s:

Turbo button: This button, if you have it, does nothing. It's a holdover from the early days of the PC when the computer ran in two modes: very slow and as fast as it could. Obviously, most people would opt to run their computers fast because that's what they paid for. The slow setting is rarely, if ever, used. A small light accompanies the Turbo button.

Keyboard lock: For about ten years, between 1984 and 1994, most PCs came with a tiny key and lock. You used the key to lock the keyboard; when the keyboard was locked, the computer ignored what you typed. Some locks even prevented you from opening the computer's case and getting inside. Mostly, these locks were for show, which is why the fad faded out several years ago.




- ✓ The console isn't the only part of your computer system that sports an on-off switch. Your PC's monitor, printer, and modem (and almost everything else) also have their own on-off switch. See Chapter 4 for more information about turning everything on.
- ✓ The on-off symbol shown above may indicate the Reset button on some computers. Check with your computer manual to be sure.
- ✓ Try not to block the air vents on the front of the console. If you do, the computer may literally suffocate. (Actually, it gets too hot.)
- ✓ Look in Chapter 11 for information about the horrors that lurk on the console's ugly backside.
- ✓ If your computer does have a lock and key, don't count on using it for security purposes: I have several computers in the office from different manufacturers, and the same key works with all the locks.
- ✓ Hard drive lights can be red or green or yellow, and the light flickers when the hard drive is in use. Don't let it freak you out! It's not an alarm; the hard drive is just doing its job. (Personally, I find the green type of hard drive light most comforting — reminds me of Christmas.)

Stuff on the console (back)

As computer designers strive to make their product prettier, they've moved many of the important connections and doodads to the PC's rump. Figure 1-3 shows you where some of the doodads are located and what they connect to. Your computer probably will have most of the items shown in the figure, although they'll probably be in a different location on the PC's backside.

Power connector: This thing is where the PC plugs into a cord that plugs into the wall.



Keyboard connector: The keyboard plugs into this little hole. On some very old PCs, the hole is much larger.



Mouse connector: It's generally the same size and shape as the keyboard connector, although this hole has a mouse icon nearby to let you know that the mouse plugs in there.



USB port: Plug snazzy USB devices into these Certs-size slots. More about what can be plugged into a USB port can be found in Chapter 11.



Serial, or COM, ports: Most PCs have two of these, labeled COM1 and COM2. It's where an external modem or sometimes a mouse is plugged in. New PCs have 9-pin serial ports, and older PCs may have 25-pin ports.



Printer port: The PC's printer plugs into this connector.




Figure 1-3:
Important
doodads on
the back of
the console.



Joystick port: This port is used mainly for scientific applications.



Monitor connector: Your PC's monitor plugs into this hole. Sometimes the hole is on an expansion slot and is unlabeled. If so, you can tell what the monitor connector is because it has 15 little holes in it — more than the serial port, which is the same size and has only 9 holes.



Speaker/sound-out jack: It's where you plug in your PC's external speakers, or where you would hook up the PC to a sound system. (The USB port can also be used for external speakers.)



Line-in jack: This jack is where you plug in your stereo or VCR to the PC for capturing sound.



Microphone jack: The computer's microphone plugs into this hole.

In addition to the ports, jacks, and holes on the back of the console are expansion slots. They're the backsides of various expansion cards you plug into your PC. Some expansion slots have connectors for other PC goodies as well.

The good news? All this stuff is connected only once. Then your PC's butt faces the wall for the rest of its life, and you never have to look at it again.

"So where is my A drive?"

The A drive is your computer's first floppy drive. It's the only floppy drive if you have one, and it's typically the *top* floppy drive if you have two.

Then again, it could be the *bottom* floppy drive.



To find out which floppy drive is which, watch them when your PC starts up. The first floppy drive, drive A, has a light on it that lights up for a few moments after the PC starts. Immediately write *Drive A* on that drive by using an indelible marker, or use a label maker to create a label for the drive.

The Cheat Sheet inside the front cover of this book includes space for you to jot down your drive A location. That way, you always remember it.

- ✓ Your first hard drive is always drive C. If you have a second hard drive, it's drive D.
- ✓ Chapter 7 describes all this disk-drive-lettering nonsense in crystal-clear detail.



(Become a master of hieroglyphics)

Along with all the lights and switches, the typical computer console sports a whole Nile full of symbols. No one ever tells you what they are because they're supposedly international symbols (and even aliens from space would be able to discern their functions without consulting an intergalactic dictionary). In any event, I've listed them all for you in Figure 1-4 in case you stumble over one you cannot recognize and an alien from space isn't handy.

- ✓ Forget seeing On or Off on a computer switch. To be more politically correct, computers use a bar for On and a circle for Off (as shown in Figure 1-4). You can remember which is which by keeping in mind that a circle is an O and the word *off* starts with the letter O. (Then again, so does *on*. Just don't think about it.)

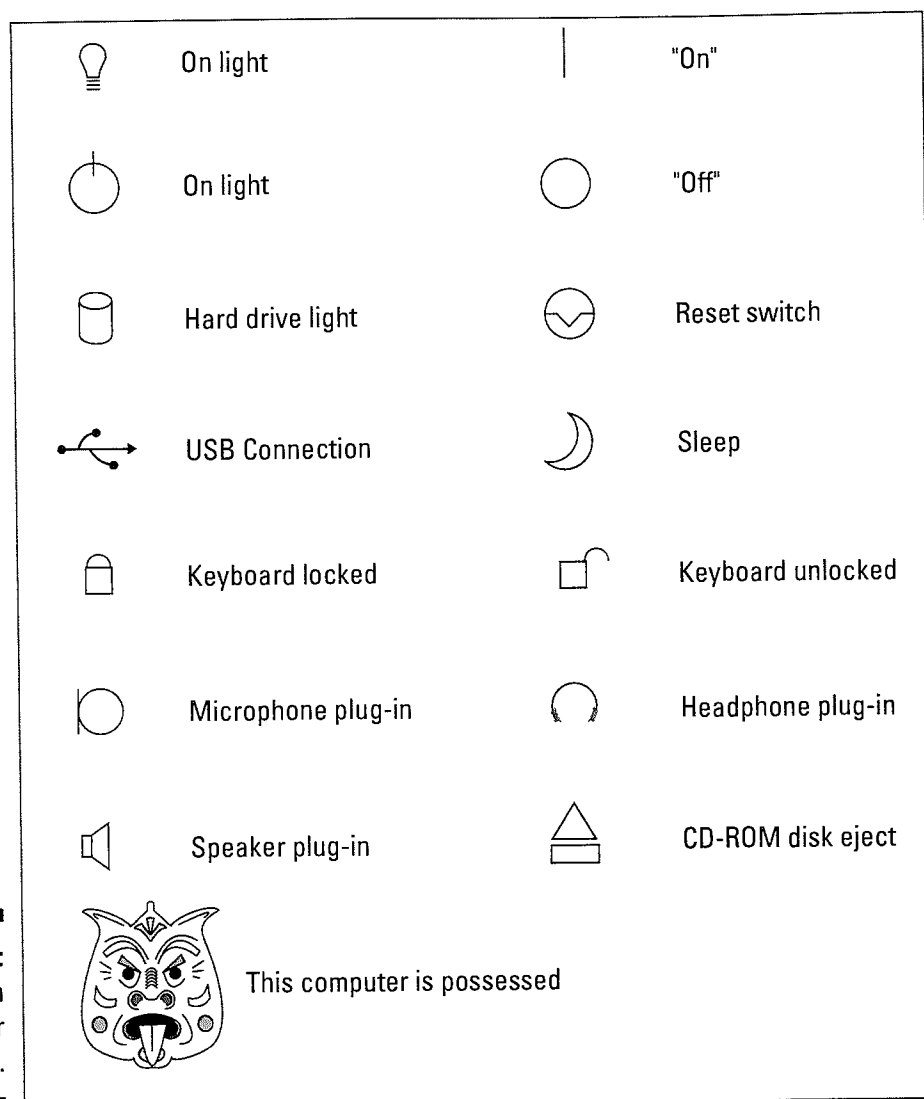


Figure 1-4:
Common
computer
hieroglyphics.

- ✓ To drive this confusing point home: Most PCs now have a dual On-Off switch with *both* symbols on it. Press once to turn on, and press again to turn off.
- ✓ Most consoles have a little light that lets you know that the computer is on. This little light is accompanied by a special symbol. In Figure 1-4, three "the computer is on" symbols are shown. What? Would it be that difficult to beat the word *On* into a foreigner's brains? In any event, when the computer is on, it makes noise. That's a definite way to know.

Variations on the typical computer theme

Not all computers look like the image shown in Figure 1-1. In fact, that's an old IBM PC shown there. Today's models are loosely based on the later IBM AT design, which is being replaced by anything slab-like that looks sleek

and has blinking lights — two of the highest status symbols a personal computer can attain. Here are some other terms used to describe various PC modes and models:

Desktop: A typical PC configuration with a slab-like console and a monitor holding everything down like a \$500 paperweight.

Desktop (small footprint): A PC's *footprint* is the amount of desk space it uses. A small footprint desktop model is just tinier than the full-size desktop model. Of course, in the end it makes no difference: The amount of clutter you have always expands to fill available desk space.

Laptops: A specialty type of computer that folds into a handy lightweight package, ideal for toting around. Laptop PCs work just like their desktop brethren; any exceptions are noted throughout this book.

Towers: Essentially a console standing on its side, making it tall, like a tower. These PCs have more room inside for expansion. They typically sit on the floor, and the monitor and keyboard are on top of the desk. Preferred by power users.

Minitowers: A small, squat version of the tower PC designed for people who work near airports where height restrictions are in effect. Seriously, you can put a minitower on top of your desk, typically next to the monitor and keyboard. A typical minitower is illustrated in Figure 1-2.

Tiffany Towers: Actually the name of a stripper and has nothing to do with computers.

Also see Chapter 31, which presents a list of ten things that may or may not be a PC.

Your Basic Software

Computer software doesn't get the credit it deserves for running your computer. That's probably why it's overpriced. In any event, you need the software to make your hardware go.

The operating system (or "Who's in charge here?")

The most important piece of software is the *operating system*. It's the computer's number-one program — the head honcho, the big cheese, Mr. In Charge, Fearless Leader, da King.

The operating system rules the computer's roost, controlling all the individual pieces and making sure that everything gets along well. It's the actual brains of the operation, telling the nitwitted hardware what to do next. The operating system also controls applications software (see the following section). Each of those programs must bend a knee and take a loyalty oath to the operating system.



- ✓ The computer's most important piece of software is the operating system.
- ✓ The operating system typically comes with the computer when you buy it. You never need to add a second operating system, although operating systems do get updated and improved from time to time. See Chapter 20 for information about upgrading the operating system.
- ✓ It used to be, in the olden days (about 1986), that you bought a program for a specific type of computer. The software store had sections for IBM, Apple, and Commodore. Now you buy software for a specific operating system: Windows, Linux, or Macintosh.
- ✓ For the PC, the most popular operating system used to be DOS. Now it's Windows. Although other popular operating systems exist, Windows is pretty much king of the heap. Bill Gates knows this fact every time he gets his monthly bank statement.
- ✓ Chapter 5 chitty-chats about Windows.

Other types of programs

The operating system is merely in charge of the computer. By itself, an operating system doesn't really do anything for you. Instead, to get work done, you need an application program. *Application programs* are the programs that do the work. They include word processors, spreadsheets, and databases. Whatever it is you do on your computer, it's being done by an application program.

Other types of programs include utilities and games and educational and programming software. Other categories might exist, but I'm too lazy right now to think of them.

- ✓ Part V of this book covers computer software.
- ✓ *Internet browsers* are programs you run on your PC to "browse" through the millions of pages that make up the World Wide Web. See Part VI for more Internet stuff.
- ✓ *Utilities* are programs that carry out special tasks, typically enhancing the capabilities of the operating system.



- ✓ Games. Well. What more can I say?
- ✓ Educational software doesn't mean only programs to teach Tommy to count. I heartily recommend typing-tutor software to teach you to be a better typist at the computer. I've even used musical software to train my ear so that I can be a better musician. (Hasn't helped much.)
- ✓ You don't have to learn how to program the computer to use it.

Even so, if you *really* want to tell your computer what to do with itself, consider picking up a programming package. One of the easiest to learn is Microsoft Visual Basic. One of the most popular is the C language. Appropriate ...*For Dummies* books on these topics are available (from IDG Books Worldwide, Inc.) if you truly want to be a master of the dopey machine before you.

One Final, Consoling Word of Advice

The last thing you should be concerned about is that your PC — your personal computer — will blow up. It'll never happen. No sparks. No flash. No boom.

In many science fiction movies, computers blow up and spew fire and rocks. Irwin Allen did that in all his 1960s TV shows. Even *Star Trek's* Mr. Spock was fond of pointing at some alien computer and uttering, in his calm Vulcan way, "Push this button, and the entire planet will become molten rubble." But in reality, it won't happen. Computers are just too dull. Sorry.

Chapter 5

The Operating System (Or “It Does Windows!”)

In This Chapter

- ▶ Exploring the reasons that computers have operating systems
 - ▶ Understanding Windows
 - ▶ Using the taskbar
 - ▶ Manipulating windows on the screen
 - ▶ Working a dialog box
 - ▶ Getting help
-

The PC’s original operating system was DOS. But DOS was ugly, and everyone complained about it, so Microsoft gave us Windows, which is pretty, and everyone complains about it. Still, your PC needs an operating system, and Windows is probably the one you’re stuck with. I know that because I’ve seen Bill Gates’s planner, and it says, in Step 9 in the section on taking over the world, “Get everyone to use Windows.”

This book specifically covers Windows 98, the latest PC operating system. Even so, because Windows 98 is only a minor update from the preceding version, Windows 95, almost all the information applies to both versions. The following tidbits of text help get you oriented to your PC’s operating system, whether you have it or not, whether you like it or not.

“Why the Heck Do I Need an Operating System?”

I’ve always figured that operating systems were unnecessary. In fact, the first thing I did when I got my first DOS computer was type `ERASE DOS`, just to see what would happen.

Nothing happened.

Well, I got a File not found error. Luckily for me, DOS didn't get erased.

Windows, like DOS, is an operating system. An operating system is necessary to run your computer. It's the software that controls everything — all the hardware. Additionally, your PC's operating system is what dishes up applications programs for you — serving them to you like a waiter in a restaurant. You choose Word Processing from a menu, and the operating system runs that program. Simple. Maybe even fun.

- ✓ Your primary duty with Windows is to tell it to run your software. I cover this task in Chapter 6.
- ✓ As your secondary duty, you use Windows to manage the many files and documents you create. That's another aspect of an operating system: organizing all your computer junk and storing it properly on the hard drive. Chapter 8 covers this subject.
- ✓ Your tertiary (meaning *third*) duty in Windows is to run your computer. It's the geeky aspect, the thing that drives too many people over the edge. Might be covered in this book. Might not. I haven't made up my mind.

Windows, Your PC's Real Brain

The main program in charge of your PC is Windows. Ideally (which means that it could never happen in real life), a computer's operating system should be quiet and efficient, never getting in the way and carrying out your instructions like a dutiful and grateful servant.

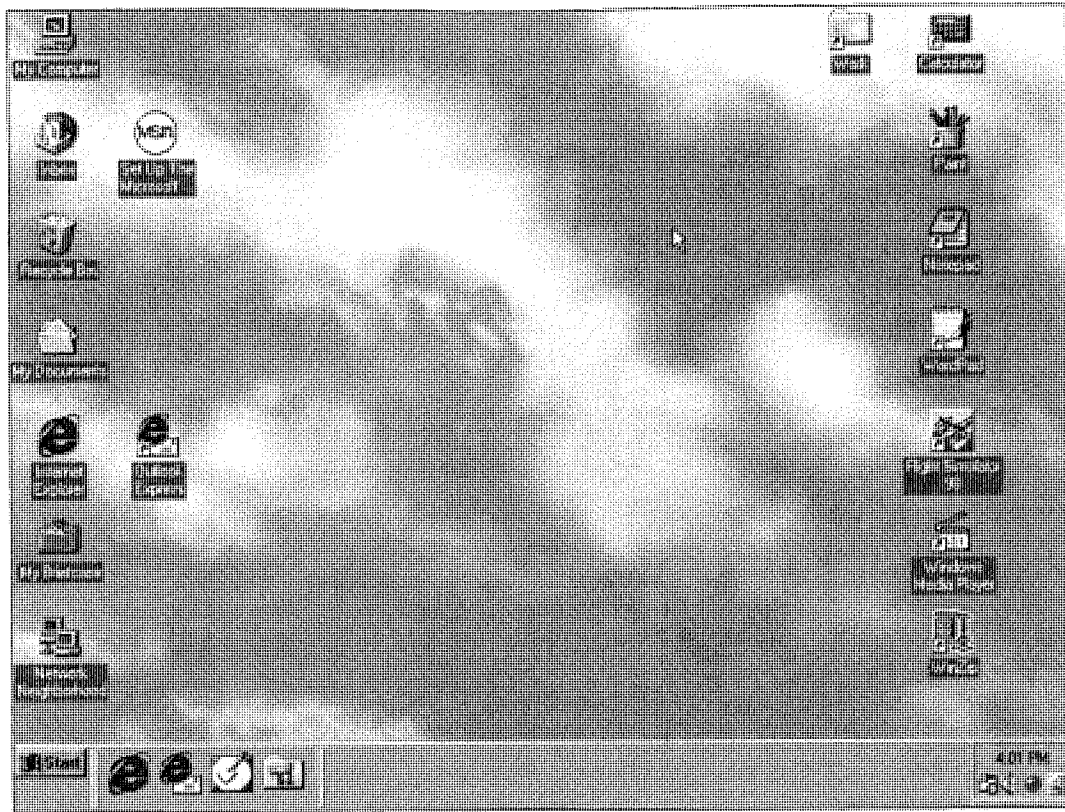
In reality, Windows is a rude little kid. It behaves like an arrogant teenager who's handsome and fun but who won't tell you where he has hidden your car keys or your wallet unless you play poker with him. In other words, with Windows in charge of your PC, you must play the game by Windows rules.

Where is the desktop?

Windows works graphically. It shows you graphical images, or *icons*, representing everything inside your computer. These graphics are all pasted on a background called the *desktop*. In Figure 5-1, the desktop has the famous Windows clouds background.

You control everything by using your computer's mouse. The mouse controls the pointer on the desktop, which looks like an arrow-shaped UFO in Figure 5-1. You use the mouse and its pointer to point at things, grab them, drag them around, punch 'em, scratch 'em till they bleed, and other mouse-y things like that.

Figure 5-1:
The
Windows
desktop.



Oh, you can also use the keyboard, although graphical operating systems such as Windows love mice more than they love keyboards.

- ✓ The *desktop* is merely the background on which Windows shows you its stuff — like an old sheet you hang from the wall to bore your neighbors with your Cayman Islands vacation slide show.
- ✓ The little pictures are called *icons*.
- ✓ Figure 5-1 shows what Windows may look like. On your computer, it looks different (probably because your computer doesn't like you).
- ✓ Refer to the end of Chapter 14 for more information about using a mouse, including all those mouse activities and their associated terms.
- ✓ Using your keyboard is covered somewhat in Chapter 15.

Between labors, Hercules did not go to the taskbar

That gunboat-gray strip along the bottom of the desktop is called the *taskbar*. It's the Windows main control center.

On the left end of the taskbar is the Start button. Yes, that's where you start programs in Windows. You can also shut down Windows by using the Start button. Start. Stop. Microsoft can't make up its mind.

On the right end of the taskbar is the *system tray*. I like to call it the loud time because it typically looks like a speaker shouting out the time of day. Other items may show up on the system tray (refer to Figure 5-1, for example). If you don't have a sound system in your PC, the speaker doesn't show up. If you're computing on the international date line, the time doesn't show up either.

From time to time, buttons appear in the middle of the taskbar. Each button represents a window or program you have floating open on the desktop. Or it can represent a program you've put away or *minimized*, which I cover in Chapter 6. All this means something, which I probably get into later.

The taskbar can also be home to various toolbars in Windows 98. For example, Figure 5-1 shows the Quick Launch bar just this side of the Start button. That toolbar contains buttons that let you quickly start programs, and it may or may not appear on-screen, depending on the number of sunspots this month.

- ✓ You can point the mouse at the various items on the system tray to get more information or to control them. Clicking the items usually does something, depending on what and how you click. For example, double-click the time, and you can set the computer's clock (it's covered in Chapter 11).
- ✓ The taskbar can float on any edge of the desktop; use your mouse to drag the taskbar to the top, left, or right sides of the screen. (Point the mouse at a blank part of the taskbar to drag it.) Most folks leave it on the bottom, which is where this book assumes that it lies.



The almighty Start button

Everything in Windows starts with the Start button, conveniently located on the left side of the taskbar. The Start button controls a pop-up menu (and submenus galore!), on which you find various commands and programs.

To pop up the Start menu, click its button by using your mouse. Click.