CENTER FOR DEMOCRACY & TECHNOLOGY, AMERICAN CIVIL LIBERTIES UNION; PLANTAGENET, INC
v.
GERALD J. PAPPERT, Attorney General of the Commonwealth of Pennsylvania

337 F. Supp. 2d 606 (E.D. Pa. 2004)

September 10, 2004, Decided

III. FINDINGS OF FACT

* * *

B. INTERNET OVERVIEW

1. Technical Overview of the Internet and the World Wide Web

6. The Internet is a global "network of networks" that allows Internet users to send and receive a huge diversity of content and communications. The "World Wide Web" is a common method that Internet users can use to make content available to other Internet users. Jt. Stip. P6.

7. In the United States, most people access the Internet through ISPs. Home Internet users generally contract on a monthly or annual basis with an ISP and will access that ISP's network over a dial-up telephone line, or a higher-speed connection such as a cable or digital subscriber line ("DSL"). A typical ISP's network is in turn connected, directly or indirectly (through a larger ISP), to the network of an Internet backbone [**13] provider (a very large ISP with high-speed transcontinental or global data lines), and through the backbone to other backbones, ISPs, and networks that, collectively, comprise the global Internet. Jt. Stip. P7.

8. Businesses in the United States commonly contract with an ISP to provide Internet access to their employees or to connect their internal computer network to the ISP's network, which is in turn connected to the global Internet. Some businesses connect to their ISP's networks (and the Internet) over dedicated high-speed connections, while other businesses access the Internet over dial-up telephone lines, cable circuits, or DSL circuits. Jt. Stip. P8.

9. A communication over the Internet will commonly travel up the "tree" or hierarchy of networks of one or more backbone providers and then back down to its destination. A hypothetical communication (from an employee of a corporation) might originate on the user's computer, travel through the corporation's network, then through a regional ISP's network, then to a backbone provider, then to another backbone provider, then back down to a regional ISP, then, in some cases, through the network of a smaller ISP, and then to the corporate [**14] network of the destination, and finally to the computer of the intended recipient of the communication. Pls.' FOF P34; Tr. 1/6/04 (Marcus) pp. 66-67; Pls.' Ex. 2 (Marcus Expert Presentation) at 1.

10. The Act is the first attempt by a state to impose liability on an ISP that does not directly contract with the originator of a communication. Thus, Pennsylvania is the first state to impose liability on, [*614] for example, the backbone providers or regional ISPs that route the communication in the hypothetical example in Finding of Fact 9.

11. Some communications on the Internet are divided into small "packets" that are separately sent over the Internet and reassembled on the receiving end. Pls.' FOF P 36; Tr. 1/6/04 (Marcus) pp. 68-69; Tr. 2/26/04 (Marcus) pp. 33-34. Separate packets that make up a given communication on the Internet are not required to travel over the same path from the sender to the recipient of the communication but can be routed over different paths within an ISP's network, or in the middle of the Internet, based on a variety of factors such as congestion on the network. Pls.' FOF P 38; Tr. 1/6/04 (Marcus) pp. 70-71.

2. Publishing to the World Wide Web

[**15] 12. Individuals, businesses, governments, and other institutions that want to make content broadly available over the Internet (hereafter "web publishers") can do so by creating a web site on the World Wide Web. Jt. Stip. P9.

13. To make a web site available on the World Wide Web, a web publisher must place the content or web pages onto a computer running specialized web server software. This computer, known as a Web Server, transmits the requested web pages in response to requests sent by users on the Internet. Jt. Stip. P10.

14. Web publishers have two common options for making a web site available over a Web Server. First, a web publisher can own and operate a Web Server on the web publisher's premises (including, possibly, the web publisher's home). In that case, the web publisher would contract with an ISP for Internet access and would thereby connect the Web Server to the Internet. Jt. Stip. P11.

15. Second, a web publisher may contract with a web host (or web hosting company) to own and operate the necessary Web Server on the web host's premises (or third party premises arranged by the web host). A web host will typically operate one or more Web Servers that can store [**16] web pages for customers and make those web pages generally available to users on the Internet. Many ISPs offer web hosting services, but many web hosts operate independently of ISPs. Jt. Stip. P12.

16. A web host offers a web publisher the ability to post a web page or a web site to the World Wide Web. There are a variety of forms of web hosting, including arrangements where a web hosting company: (1) provides a Web Server to service a single web site of a customer, (2) provides a Web Server the customer can use to run multiple web sites, or (3) provides space on a Web Server that services the web sites of many different customers. The third form of web hosting is commonly called virtual web hosting. Pls.' FOF PP 84, 86; Tr. 1/7/04 (Clark) pp. 178-81.

17. To access web pages on a web site, an Internet user utilizes a client computer program called a web browser. Microsoft Explorer and Netscape are two common web browsers. A web browser sends a request to a Web Server, which responds by sending the requested web page,

which upon receipt is formatted and displayed by the web browser. Pls.' FOF P 43; Tr. 1/6/04 (Marcus) p. 72-73.

3. Domain Names and URLs

18. Typically, [**17] but not always, when creating a web site, a web publisher obtains a domain name that can be used to designate and locate the web site. For example, defendant obtained the domain name "attorneygeneral.gov" for his web site. Jt. Stip. P13.

19. At a minium, a domain name contains a top level domain name and a second [*615] level domain name. The following are some top level domains available in the United States: .com, .net, .org, .edu, .gov, .biz, .info. Most nations in the world also have country top level domains, such as us (United States), uk (United Kingdom), es (Spain), it (Italy), ru (Russian Federation). The second level domain appears to the left of the top level domain, separated by a dot. One acquires a domain name (top level and second level) by purchasing it from, and registering it with, a registrar designated for the relevant top level domain, or, alternatively, by purchasing it from someone who already owns the name and offers it for sale. The owner of a domain may create sub-domains that are identified to the left of the second level domain and separated from it by a dot, e.g., subdomain.attorneygeneral.gov. Tr. 1/29/04 (Stern) pp. 32-33; 2/18/04 (Stern) pp. 58, 106-108; [**18] 3/1/04 (Stern) p. 106; Tr. 1/6/04 (Marcus) p. 76; Tr. 1/7/04 (Clark) pp. 137, 151, 154, 156; Tr. 3/1/04 (Blaze) p. 58.

20. Domain names are read right to left. The part of a domain name furthest to the right (or the top level domain) is the broadest part of the domain name. As the domain name is read to the left, the sub-domains identify specific Web Servers or web sites. For example, "upenn.edu" is a sub-domain of the .edu top level domain identifying the University of Pennsylvania's Web Server, and "cis.upenn.edu" is in turn a sub-domain of upenn.edu used to identify a web site for the Department of Computer and Information Science at the University of Pennsylvania. Pls.' FOF P45; Tr. 1/6/04 (Marcus) p.76.

21. A domain name can be coupled with additional information to create a Uniform Resource Locator ("URL") which is a more complete way to designate certain content or other resources on the Internet. Jt. Stip. P14.

22. A URL is the commonly used textual designation of an Internet web site's address. Thus, for example, the URL of defendant's web site is http://www.attorneygeneral.gov. The http indicates that the Hypertext Transfer Protocol (which is the main protocol used to [**19] transmit World Wide Web pages) is to be used. The "www.attorneygeneral.gov" part of the URL is used to locate the specific Web Server(s) that contains (hosts) the content for the requested web site. Jt. Stip. P15.

23. A web page accessed by a URL like http://www.attorneygeneral.gov is commonly referred to as the home page of the web site. A URL can also contain a reference to a specific sub-page contained in a web site. The sub-page is designated in writing by slashes after the home page (such as http://www.attorneygeneral.gov/press/pr.cfm). A single web site can contain thousands of different web pages. Jt. Stip. P16.

24. A URL on the World Wide Web only refers to a location where content can be found. A URL does not refer to any specific piece of static content - the content is permanent only until it is changed by the web site's webmaster (often, but not always, the owner of the web site). The actual content to which a URL points can (and often does) easily change without the URL changing in any way. Pls.' FOF P 53; Tr. 1/6/04 (Marcus) p. 77; Tr. 1/6/04 (Blain) pp. 26-28.

4. Browsing the Web

25. For accessing content on the World Wide Web, the most common sequence [**20] is for a user to request content from a web site, and for the web site to return web pages to the user. This sequence is illustrated as follows, with the initial request shown by the arrows on the left, and the response shown by the arrows on the right:

[*616] User

-

User's ISP

-

Internet "Backbone" Provider(s)

-

Web Site's ISP

-

Web Site

Jt. Stip. P22.

26. In the vast majority of cases, the user's ISP is different from the web site's ISP. Jt. Stip. P23.

27. To access a web page, a user can either type the URL of the web page into his web browser, or, if the user is already accessing a web page, click on a hyperlink that takes the user to a different web page. A hyperlink is commonly shown on a web page with underlining; for example, on a web listing of the University of Pennsylvania Department of Computer and Information Science faculty members, an individual professor's name would be underlined and clicking on the name would take the user to the professor's personal web page. Pls.' FOF P57; Tr. 1/6/04 (Marcus) pp. 78-79.

5. Shared Domain Names

28. Within the United States alone, there are tens of millions of separate domain names used for web sites [**21] that are, for the most part, independent of each other. In the great majority of those situations, a single web publisher controls the domain name and the entire web site and is responsible for all pages and sub-pages on a web site. Thus, www.example.com could be the

fully qualified domain name for a single web site (with multiple pages) controlled, hypothetically, by the Example Corporation. This approach - of a single web site being coextensive with the domain name - is the most familiar approach to placing content on the web. Jt. Stip. P17.

29. Web publishers can also publish on the World Wide Web without obtaining their own unique domain names for their web sites. For example, a web publisher can place content with a provider that offers to host web pages on the provider's own web site (as a sub-page under the provider's domain name). Thus, hypothetically, the Example Corporation could have a web site at the URL http://www.webhostingcompany.com/example. Some such providers offer their users discussion forums, chatrooms, and other services and are known more broadly as online communities. Jt. Stip. P18.

30. GeoCities is an example of an online community located in the United States. [**22] GeoCities hosts web pages of its users as sub-pages of its domain name. As an illustration, the Association of Black Women Lawyers of New Jersey, Inc., is part of the GeoCities online community, and its web pages are available at the URL http://www.geocities.com/abwlnj/homepage.html. Jt. Stip. P19.

* * *

32. Outside of the United States, www.terra.es is a well-known Spanish-language online community providing web hosting services. Jt. Stip. P20.

33. Some web hosts allow users to create web sites using individualized subdomains of the web hosts' primary domain. Thus, hypothetically, the Example Corporation web site might be at the URL http://example.webhostingcompany.com, while another customer site might be at [*617] the URL http://acehardware.webhostingcompany.com. Jt. Stip. P21.

34. Other than their existence as sub-pages or sub-domains [**23] on a providers' domains, web sites hosted as sub-pages or sub-domains are usually independent of the provider and independent of each other. Tr. 1/29/04 (Stern) pp. 54-59; 2/18/04 (Stern) pp. 103, 104.

35. Many web hosting companies offer to host web sites at a very low cost and often theses hosts offer virtual hosting - they host sites on a single Web Server. Some such web hosting companies offer to host web sites at no charge in exchange for the right to place advertisements for their service on the customer's web site. Pls.' FOF P96, Tr. 1/6/04 (Blain) pp. 26-30 (describing creation of free and low-cost web sites for two community organizations); Tr. 1/7/04 (Smallacombe) pp. 81-83, 100-102.

6. IP Addresses and the Domain Name System

36. A URL such as http://www.attorneygeneral.gov or http://www.geocities.com/abwlnj/homepage.html provides enough information for a user to access the desired web site. The user enters the URL in her web browser. However, the URL alone is not sufficient for the user's computer to locate the web site. A user's computer must first determine the numeric Internet Protocol Address or IP address of the desired web site. Every device, or [**24] computer, using the Internet must have a unique IP address.

37. When a user seeks to access a particular URL, the user's computer initiates a look up through a series of global databases known as the domain name system ("DNS") to determine the IP Address of the Web Server that can provide the desired web pages. Jt. Stip. P 24, Tr. 1/29/04 (Stern) p. 36-40. To search for the requested URL's IP address, the user's web browser must query a domain name system server ("DNS server") that has been assigned or selected within the user's computer. That DNS server attempts to find the IP address of the fully qualified domain name specified in the URL entered by first looking in its own database of domain name/IP address combinations. If that DNS server cannot find the IP address in its own database, it queries other DNS servers until it receives the correct IP address. It then returns that address to the user's computer. This process is referred to as resolving a hostname to its IP address. Jt. Stip. P25.

38. Typically, an ISP gives its customers the IP addresses of DNS servers controlled by the ISP. The addresses are entered in the customers' computers during the Internet access set-up process, [**25] a process that is often automated. Some ISPs assign a new IP address identifying a different DNS server each time the user establishes a connection to the ISP. This is called dynamic assignment. Tr. 1/29/04 (Stern) p. 36; Tr. 1/7/04 (Marcus) pp. 5-7; Tr. 1/7/04 (Smallacombe) pp. 113-114; Tr. 1/27/04 (MacDonald) pp. 144-148.

39. Companies and other network operators can choose to operate their own DNS servers. Pls.' FOF P 65, Tr. 1/6/04 (Marcus) p. 84. Individuals can also chose not to use the DNS server assigned by their ISP and can either use a DNS server available on the Internet or operate their own DNS server. Tr. 1/6/04 (Marcus) pp. 83-84, 116.

40. The numeric IP address of the DNS server provides the user's computer with the Internet address of the Web Server to which the user's computer then sends a request for the particular URL entered in the user's web browser. IP addresses (in the most common current form) are generally expressed as four sets [*618] of numbers separated by periods, e.g., 207.102.198.176. Jt. Stip. P26.

41. IP addresses are assigned by several registries covering various parts of the world. Tr. 1/29/04 (Stern) p. 37. The party to whom the registry assigns an [**26] IP address may subassign the address. The sub-assignment may, but need not, be recorded with the registrar. The sub-assignee may further sub-assign the IP address. Tr. 1/27/04 (Krause) pp. 119-120; Tr. 1/28/04 (Clark) pp. 152-160; Tr. 1/12/04 (Guzy Jr.) pp. 36, 60-62, 181, 182.

42. Although a specific URL refers only to one specific web site, many different web sites (each with different domain names and URLs) are hosted on the same physical Web Server, and all the web sites on a server share the same IP Address. Jt. Stip. P27.

43. It is common for web hosting companies to offer virtual web hosting, discussed in Finding of Fact 16, under which many web sites are hosted on the same Web Server and thus share the same IP address. Pls' FOF P 72; Tr. 1/6/04 (Marcus) p. 94; Pls.' FOF P73; Tr. 1/29/04 (Stern) p. 65; Dep. of G. Lipscomb (Comcast) at 115-16; Dep. of C. Silliman (WorldCom) at 103 (anecdotally from general industry information, it is believed that the majority of web sites share IP addresses with more than 50 sites). As an example of virtual web hosting, PlantageNet hosts about 160 to 170 of its web hosting customers - all with their own unique domain names - on a

single Web [**27] Server with a single IP address. Tr. 1/7/04 (Smallacombe) pp. 82-83. As another example, discussed more fully below, Laura Blain's web site shared its IP address with more than 15,000 other domains. Tr. 1/7/04 (Clark) pp. 141-42.

44. Research by plaintiffs' expert Michael Clark empirically confirms the prevalence of shared IP addresses. In October-November 2003, Mr. Clark created a database of 29.5 million domain names and the IP addresses to which each domain named resolved. Using this database, which was received in evidence in CD-ROM form as Plaintiffs' Exhibit 77, Mr. Clark analyzed the frequency with which IP addresses were shared among domain names. Pls' FOF PP 76. Tr. 1/7/04 (Clark) pp. 134-35, 137-40, 151-60, 170-71. In Joint Stipulation 59, the parties agreed that, for a variety of reasons, it is difficult to state a precise percentage of domain names that share an IP address with other domain names. However, they agreed and stipulated that "at the time the data was collected (October 2003), at least fifty percent of domains shared an IP address with at least fifty other domains." Pls.' FOF P 78; Jt. Stip. P 59. Some domains do not share IP addresses with other domains but [**28] are the only domain located at a single IP address. As of October 2003, over 2.5 million domains had their own, unshared, IP addresses. Jt. Stip. P59.

45. When a request for a web site reaches a Web Server that supports multiple web sites, the Web Server reads the request, including the IP address and the URL, in order to determine which web site is being requested, and returns only the requested web page or other resource. Jt. Stip. P28.

46. When a request for a particular web page is sent by the user's web browser to a Web Server, no ISP that carries the request must read the details of the request - an ISP routing the request is only required to read the destination IP address, and the ISP would effectively not be aware of the specific web site or URL requested. Pls.' FOF P 70; Tr. 1/6/04 (Marcus) pp. 92-93.

47. One cannot determine with any certainty - using technical means - whether a given web site shares its IP address with another web site. The most reliable method [*619] of determining whether a particular web site uses an IP address shared by other web sites is to contact the web hosting entity. Tr. 1/7/04 (Clark) pp.182-83. As Mark Krause, Senior Manager of Internet Infrastructure [**29] Security for WorldCom/MCI, explained, it is "hard, or impossible for [an ISP] to determine what other content" might be behind a particular IP address. Pls.' FOF P 80; Tr. 1/27/04 (Krause) pp. 9, 98.