

specialized knowledge, education, and expertise as applied to the facts and circumstances in this case.

Background

3. Hotfile is a website that allows its users both to “upload” electronic files and to “download” files uploaded by others. Hotfile provides any user who uploads a file with a “link” (an alphanumeric string that operates as a uniform resource locator or “url”). The uploaded file can then be downloaded by any Internet user who has access to the link. Users who download from Hotfile find “links” to files hosted on Hotfile in a variety of ways on the Internet (such as on other websites that host collections of such links, or via email or an online forum). They can then follow the link to Hotfile and download from Hotfile the file referenced by the link.

4. Entertainment content, such as music, motion pictures, television shows, and so forth, can be represented as a digital computer file, and can easily be communicated over the Internet in digital form. Thus, it is easy for users to upload and download entertainment content (such as motion pictures and television shows) to and from Hotfile, as described above.

Information Maintained by Hotfile

5. I have not yet had the opportunity to examine Hotfile’s back-end computer systems; I have reviewed and analyzed what is accessible and visible to the public. Nonetheless, based on the features of the website, it is possible to draw certain conclusions about the types of data that the site must necessarily be using and/or maintaining. I describe these features below.

6. First, Hotfile must maintain copies of files uploaded by users and available for downloading through the site. For purposes of this Declaration, I will refer to these files as “Content Files.” The effective functioning of Hotfile as a content dissemination site requires that it be possible, by examining these files, to determine the content embodied therein (for example,

if a Content File represented a particular motion picture, that fact could be determined by examining the file itself).

7. Second, Hotfile must maintain certain information about the Content Files. As described below, certain features of Hotfile require it to use certain information about the Content Files on its system. It is common for computer systems to maintain a database (or databases) or index of files that the computer system stores, separate and apart from the files themselves. Just as a personal computer user can use “Windows Explorer” or “Mac Finder” to list the files on their computer, Hotfile must maintain information about the Content Files on its system. Another analogy is a card catalog at a library, which provides information about the books in the library, maintained separately from the books themselves. For purposes of this Declaration, I will refer to this information about Content Files as “Content Reference Data.”

8. I expect the Content Reference Data to be separate from, and orders of magnitude smaller than, the Content Files. Indices or databases are conventionally much smaller than the files they index, just as a card catalog is much smaller than the library. Even if the Content Files themselves are voluminous, it should not be a particularly difficult task to generate a copy of the Content Reference Data. The Content Reference Data, in turn, would provide a complete list of the Content Files available on the Hotfile system, as well as additional information about those files as described below. Access to this list would make it possible to request specific Content Files from Hotfile.

9. Third, Hotfile must maintain certain information about its users. As described below, Hotfile allows any user to upload and download files, but also allows users to become “premium” users (in exchange for a fee); premium users enjoy certain privileges (such as the ability to download files more quickly and conveniently). Hotfile also compensates users under

a variety of conditions (for example, users who upload files are compensated based in part on how often their files are downloaded). In order for these features to work, Hotfile must maintain certain categories of information about users. For purposes of this Declaration, I will refer to this as “User Data.” As with the Content Reference Data described above, I would expect the User Data to be orders of magnitude smaller than the Content Files themselves, and it should not be difficult to generate a copy of such data. Because Hotfile would need to use such User Data as part of its regular operation, I expect this data to be consolidated in a location from which it could be readily copied.

10. Hotfile must also have access to certain information about downloads of files stored on its system. It is common practice for servers used to deliver digital content to make a record of each incoming request and outgoing response. There are a variety of reasons that web servers generate this information; it can be used for technical reasons (such as troubleshooting customer problems) and/or for business reasons (such as analyzing website load and to determine whether additional web servers are required). In addition, website operators will typically maintain other records relating to their specific business functions (such as information related to the billing of their customers). For purposes of this Declaration, I will refer to these kinds of information as “Activity Data.” As described below, Hotfile must be utilizing and maintaining at least some of this information in order to operate certain features of its website.

Information Maintained in Content Reference Data and User Data

11. As described above, I have not yet had the opportunity to directly observe the Content Reference Data, User Data and Activity Data maintained by Hotfile, only the features of the site that are available and visible to the public. But based on my observation of the operation

of the Hotfile system, I can draw some conclusions regarding the types of information that Hotfile must maintain.

12. In the discussion below, I group this information into two sections – “Content Reference Data” and “User Data.” The information I have previously described as “Activity Data” will overlap with both categories, as Hotfile must link downloading activity both to certain users and to the files hosted on its system. Lacking the benefit of direct observation of Hotfile’s back-end systems, I do not yet know how Hotfile has chosen to organize the information described below. Therefore, information I list as “Content Reference Data” may in practice be grouped with “User Data,” or vice versa, depending on how Hotfile has chosen to organize its databases. Regardless of the exact organization, as noted above, this information is likely to be maintained separately from the Content File data and in a form in which it can be readily copied.

13. With respect to the Content Reference Data, I have observed that a user with a Hotfile account (a “registered user”) can view, when logged in to the site, a complete list of the Content Files they have uploaded, the link(s) associated with each Content Files they have uploaded, the size of each Content File they have uploaded, and, within certain parameters, the number of times each Content File they have uploaded has been downloaded. In order to accurately render this information, Hotfile must maintain, for each Content File, the:

- a. The link(s) associated with that file;
- b. The size of the file;
- c. The identifier associated with the user who uploaded the file; and
- d. The number of times the file has been downloaded. (According to the Hotfile website, this figure excludes duplicate downloads by premium users and downloads from countries other than those for which Hotfile “counts” downloads.).

In addition, information on the Hotfile website suggests that Hotfile must use information regarding each request to download a Content File. This information includes, at a minimum, (1) the location (at least at the country level) from which each file is downloaded, in order to determine whether each download originates from a country from which Hotfile says it “counts” downloads; and (2) whether or not the requesting user is a “Premium” user, information that Hotfile uses to determine both (a) the download speed to provide to the user and (b) whether to “count” the download.

14. With respect to the User Data, Hotfile, according to its website, operates “Affiliates” programs that compensate registered users under certain conditions. *See* <http://www.hotfile.com/affiliate.html>. First, under the basic “Affiliate” program, uploading users are paid when Content Files uploaded by the user are downloaded by other users (with the compensation generally depending on the number of such downloads and the percentage of downloading users who upgrade to “premium” subscriptions based on those downloads). Second, under the “Referral program[] for site owners,” users are paid when a different user accessing Hotfile via a “link” hosted on a website registered by the original user earns money through the “Affiliates” program. Third, under the “Refer a friend” program, users are paid when a different user referred by the original user upgrades to a premium account. In order for these “Affiliate” programs to be operating as described, Hotfile must maintain, for each registered user:

- a. Basic account information;
- b. Which Content Files the user has uploaded;
- c. How many other users upgraded to a “premium” Hotfile account while downloading a Content File the user uploaded;
- d. Which of the user’s Content Files other users have downloaded (and how many of those downloaders were “premium” users);

- e. Which Content Files the user has downloaded (if the user is a premium user);
- f. Any websites the user has registered with Hotfile's "Affiliate" program;
- g. If the user is a premium user that was referred to Hotfile under the "Refer a Friend" program, the identity of the referring user;
- h. Information associated with recent payments from the user (if the user is premium) and pending payments to the user (if the user participates in an "Affiliate" program).

In addition, for users who upgrade to "premium" accounts after being referred to Hotfile by a registered "Affiliate" website, Hotfile must use information regarding the identity of the referring website.

Information Maintained By Third Parties

15. Some information regarding visitors to the Hotfile website is likely to be maintained by Google. I can see, by examining the structure of the Hotfile webpages, that they are using a service called "Google Analytics." Google Analytics is a service that allows the owner of a website to maintain records of their website traffic at Google, and to analyze various properties of that traffic. Part of the operation of this service involves the transmission of information to Google and the storage of that information on Google's computers for analysis by the website owner. A description can be found at <http://www.google.com/analytics/>.

Role of Lemuria

16. Hotfile is served by an Internet Service Provider called "Lemuria Communications, Inc." ("Lemuria"). As far as I can tell, Lemuria provides ISP functions to only Hotfile.com's servers. Records show that Lemuria supports two "blocks" of Internet addresses (containing 3,072 Internet addresses). This can be determined by querying any number of publicly available online services that lookup and return information about the properties of an Internet Service Provider. Using a method called a "Reverse DNS Lookup" shows that these

addresses supported by Lemuria only resolve to Hotfile server names. Therefore, it appears that Lemuria does not provide services to any entities except Hotfile.

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

Executed on this 21st day of February 2011, at Chicago, IL.

A handwritten signature in black ink, appearing to read 'Ian Foster', written in a cursive style.

Ian Foster

Exhibit A

IAN FOSTER

NARRATIVE

I am a computer scientist whose work at the intersection of computing and the sciences has produced both practical technologies that have seen wide adoption and concepts and methods that have proven influential in research and education. My research interests span a range of topics in parallel, distributed, and data-intensive computing. A unifying theme is a desire to use the power of rapid communication to accelerate discovery, whether by linking people with remote computers and data, accelerating complex computational processes, or enabling distributed virtual teams. I pursue use-inspired basic research, meaning that I employ challenging practical problems to motivate and focus work on hard problems in computer science. Over the years, these practical problems have come from such fields as environmental science, economics, high-energy physics, biomedicine, and engineering. I often build sophisticated artifacts (software and distributed systems) in order to apply, evaluate, and disseminate new concepts and methods. Thus, my work frequently requires large teams of disciplinary scholars, computer scientists, and software engineers.

EDUCATION

1985 - 1988 Imperial College London, United Kingdom
PhD, Computer Science and Diploma of Imperial College

1977 - 1979 University of Canterbury Christchurch, New Zealand
B.Sc (Hons I), Computer Science

PROFESSIONAL EXPERIENCE

2006 - Argonne National Laboratory Argonne, IL
 The University of Chicago Chicago, IL

Director, Computation Institute (CI), Argonne Distinguished Fellow, Arthur Holly Compton Distinguished Service Professor of Computer Science

[Previously Asst. Scientist 1989-1992, Scientist 1992-1997, Senior Scientist 1997-2008, Math & Computer Science Division, Argonne, Assoc. Professor 1996-2000, Professor 2000-2006, Dept of Computer Science, University of Chicago]

- Lead the Computation Institute, a cross-institutional, multi-disciplinary research institute with 100 Fellows, 80 staff, many postdocs and students, and an annual budget of ~\$20M
- Conduct research in distributed, parallel, and data-intensive computing, and publish technical articles that have seen more than 50,000 citations, yielding an h index of 90*
- Pioneer new technologies that have seen wide adoption, including grid computing (technologies, infrastructures, and applications), parallel climate models (e.g., Parallel Climate Model), and parallel programming languages (e.g., Swift)
- Establish and manage major computer science and computational science projects, including Earth System Grid, Grid Physics Network, International Virtual Data Grid Laboratory, Open Science Grid, TeraGrid, Center for Enabling Distributed Petascale Science, and Center for Robust

* H-index is a commonly used measure of scientific impact. According to <http://www.cs.ucla.edu/~palsberg/h-number.html>, mine is the fourth highest of any computer scientist in 2010.

Decisions on Climate and Energy Policy

- Lead the establishment of the international Globus Alliance open source community, and the development of the Globus Toolkit, from inception to its adoption by a broad spectrum of national and international projects
- Establish the Open Grid Forum, and play a leadership role in various other national and international projects and organizations
- Serve on national and international advisory committees, including the US Ocean Observatory Initiative and UK eScience Program
- Supervise the work of research staff and graduate students
- Teach graduate and undergraduate classes in computer science

1985 - 1998 Imperial College London, U.K.

Research Associate, Department of Computing

- Conduct research in concurrent logic programming systems
- Develop programming language technology, commercialized as Strand

OTHER PROFESSIONAL EXPERIENCE

Founder, Board Member, Chief Open Source Strategist, Univa Corporation

Technical Advisory Board, EMC Corporation (2010-), IOCOM Corporation (2005-), Entropia (2000-2004)

PUBLICATIONS AND PRESENTATIONS

More than 300 article and technical reports, and six books, in distributed and parallel computing, computational science, and programming languages. More than 100 keynote talks and seminars worldwide.

The Grid: Blueprint for a New Computing Infrastructure, I. Foster and C. Kesselman (Eds), Morgan-Kaufmann, 1999 and 2003 (2nd edition).

Designing and Building Parallel Programs: Concepts and Tools for Parallel Software Engineering, I. Foster, Addison-Wesley, 1995.

SELECTED RECOGNITION

D.Sc. (Honoris Causa), CINVESTAV, Mexico, 2010; Fellow, Association for Computing Machinery, 2009; GndWorld "Industry Leadership Award," 2006; Network World's 50 Most Powerful People in Networking, 2005; D.Sc (Honoris Causa), University of Canterbury, NZ, 2005; InfoWorld Innovator, 2003, 2004, 2005; Fellow, American Association for the Advancement of Science, 2004; R&D Magazine Innovator of the Year, 2003; University of Chicago Distinguished Service Award, 2003; MIT Technology Review, one of "Ten Technologies That Will Change the World," 2003; British Computer Society Lovelace Medal, 2002; Fellow, British Computer Society, 2002; Federal Laboratory Consortium Technology Transfer Award, 2002; R&D100 "Most Promising New Technology" Award, 2002; Gordon Bell Award, 2001; Global Information Infrastructure "Next Generation" Award, 1997; Best Paper Award, 1995 Supercomputing Conference; British Computer Society Award for Technical Innovation, 1989.