IDEAFLOOD PATENT ON PERSONALIZED SUB-DOMAINS
U.S. PATENT NO. 6.687,746

Latest Date That Material Can Qualify for Prior Art: **August 6, 1999**

I. General Description

The Ideaflood patent claims a method of establishing and accessing what are often called "personalized" subdomains on the Internet. A subdomain is part of a "top level" Internet domain, such as yahoo.com, eff.org, or university.edu. University.edu is the domain for an entire school's website. An individual student at university.edu, say one named Alice, might have a personalized web directory on the university.edu site, say http://www.university.edu/~alice/. The method claimed in this patent allows Alice to set up her subdomain so that web users can access her directory by simply typing "alice.university.edu" into their browser.

The method disclosed in the patent focuses on a means of creating and accessing subdomains using what is called a "DNS wildcard." That is, rather than inserting a separate "personalized" entry into the Internet's address directory (called the "Domain Name System" or "DNS") for each new personalized subdomain (a process that is often burdensome and slow), the provider inserts a single "wildcard" entry (for example, the use of an asterix as in "*.eff.org") into the DNS system so that all subdomain inquiries are sent to the provider's general web server for further redirection to a specific user's directory instead of having the DNS server itself provide the specific Internet location of the user's directory.

Thus, if one were to type "alice.university.edu" into a browser in conjunction with this system, the browser would send a DNS inquiry for alice.university.edu to the DNS system; the DNS system would then look at its records, see the entry "*.university.edu" and send back the IP address for the general university.edu server. The browser would then go to the general university.edu server and that server, after looking up Alice's subdirectory, would send the browser the pages stored at http://www.university.edu/~alice/.

The patent also teaches the method of setting up such a system so when a new user registers, a personalized subdirectory and associated subdomain are automatically created based on her username. Thus, when Alice first registers with university.edu, www.university.edu/~alice and alice.university.edu are automatically set up.

SAMPLE FIGURE OF METHOD FROM PATENT

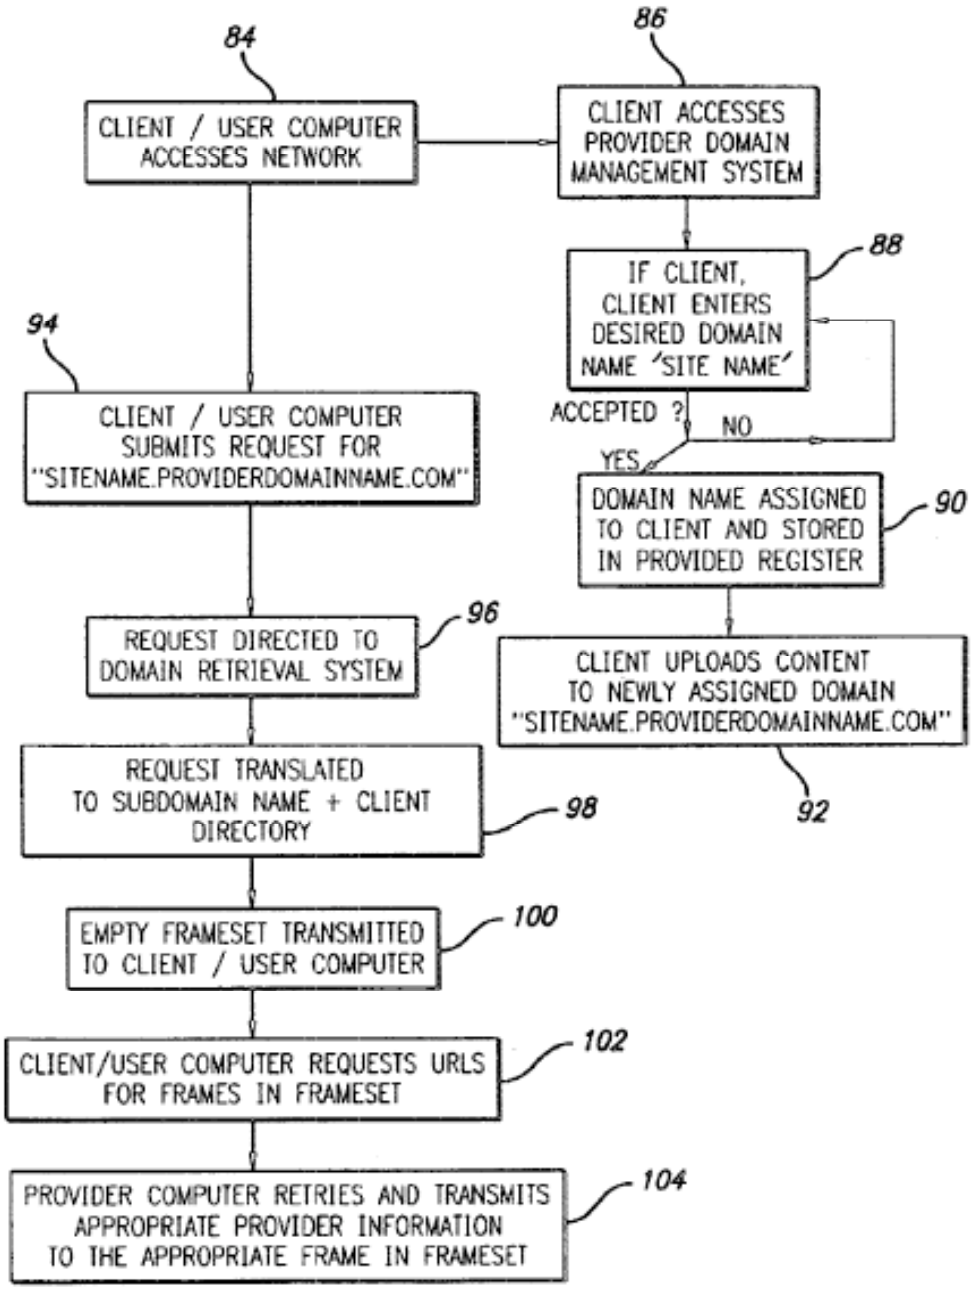
84
CLIENT / USER COMPUTER ACCESSES NETWORK
86
CLIENT ACCESSES PROVIDER DOMAIN MANAGEMENT SYSTEM
88
IF CLIENT, CLIENT ENTERS DESIRED DOMAIN NAME 'SITE NAME'
ACCEPTED ?
NO
YES
90
DOMAIN NAME ASSIGNED TO CLIENT AND STORED IN PROVIDED REGISTER
92
CLIENT UPLOADS CONTENT TO NEWLY ASSIGNED DOMAIN "SITENAME.PROVIDERDOMAINNAME.COM"
94
CLIENT / USER COMPUTER SUBMITS REQUEST FOR "SITENAME.PROVIDERDOMAINNAME.COM"
96
REQUEST DIRECTED TO DOMAIN RETRIEVAL SYSTEM
98
REQUEST TRANSLATED TO SUBDOMAIN NAME + CLIENT DIRECTORY
100
EMPTY FRAMESET TRANSMITTED TO CLIENT / USER COMPUTER
102
CLIENT/USER COMPUTER REQUESTS URLS FOR FRAMES IN FRAMESET
104
PROVIDER COMPUTER RETRIES AND TRANSMITS APPROPRIATE PROVIDER INFORMATION TO THE APPROPRIATE FRAME IN FRAMESET

II. The Claims At Issue

The patent has two independent claims and 18 dependent claims. The most important claims to eliminate are the independent claims, 1 and 12. Most of the dependent claims cover fairly obvious extensions of these that are likely to fall if the independent claims fall.

Claim 1 covers a method and claim 12 covers a system that implements the method. Claim 1 covers a method for providing Internet access to content stored at a user-selected subdomain. The method includes:

a) operating a server with an IP address stored in a DNS directory, the address consisting of a wildcard character (e.g., "*") followed by a higher-level domain name (e.g., "university.edu");
b) storing content on the server in a subdirectory with a user-specified name (e.g., alice's home directory, which the server makes available at the URL http://www.university.edu/~alice);
c) the server receiving a URL request that consists of the user-specified name followed by the higher-level domain name (e.g., alice.university.edu);
d) the server determining where the user's content is stored by extracting the user-specified name from the domain name it received (e.g. identifying www.university.edu/~alice as the location of content for alice.university.edu); and
e) the server retrieving and serving the content.

Claim One Chart

Claim Language	Example & Comments
1. A method for enabling internet access to content located by a domain name, the domain name including a user-selected subdomain label that is not associated with an IP address in a zone file of any higher-level domain, the method comprising:	Using a browser to look at the contents of alice.university.edu without having to look up the specific IP address for the subdomain in a DNS server (a.k.a. zone map).
Operating a host having an IP address specified by an internet-class resource record for a domain name server, in that the resource record associates the host IP address to a host domain name in a zone file of the domain name server, and wherein the host domain name comprises (a) a subdomain labeled with a designated wildcard character of a domain name system and (b) at least one: higher-level domain name;	Operating a website hosting server that is registered in the DNS server system with a wildcard character: "*" followed by a higher-level domain name: "university.edu".
configuring a content address according to a content storage system of the host independently of the domain name system, the content address comprising a user-selected label, wherein the user-selected label comprises at least one character that is not the designated wildcard character;	Making http://www.university.edu/~alice map to alice.university.edu through some kind of internal directory instead of using the DNS system to do it.
storing content in the content storage system, the content addressed by the content address;	storing Alice's website at www.university.edu/~alice or some other accessible location.
Receiving a domain name configured in accordance with the domain name system, the domain name comprising the host domain name with the user-selected label substituted for the designated wildcard character;	Example: http://alice.university.edu/
determining the content address from the user-selected label;	Using http://alice.university.edu to find www.university.edu/~alice, for example, using a look-up directory.
Retrieving the content from the content storage system using the content address; and serving the content.	Sending the content at www.university.edu/~alice to the requestor's web browser.

Claim 12 covers a system for providing Internet access to content identified by a user-selected subdomain name, without requiring that the subdomain name be associated with an IP address. The system includes:

a) a server whose DNS entry consists of a wildcard character followed by a higher-level domain name (e.g., "*.university.edu"); and
b) memory containing instructions that:
    1. store content on the server in a directory with a user-specified name (e.g., www.university.edu/~alice);
    2. receive a user-specified subdomain request (e.g., alice.university.edu);
    3. determine where the user's content is stored by extracting the user-specified name from the subdomain request it received
    4. retrieve and serve the content.

The dependent claims cover variations of these. Claims 2 and 13 cover the method and system of claims 1 and 12 (respectively), adding a searchable database of user-selected names that correspond to the directories as the specific method for looking up the correlation between the two. (Most of the dependent claims come in method/system pairs like this.)

Claims 3-5 and 14-15 add the steps of serving HTML frames with advertisements in one frame and the user content in another.

Claims 6-7 and 16-17 cover using schedulers and redirectors that perform some of the steps of claims 1 and 12 for load-balancing and efficiency.

Claims 8 and 18 cover receiving the content from a user who selected the user-selected name.

Claims 9 and 19 cover comparing the user-selected name to names that already exist in the system.

Claims 10 and 20 cover selling the user-selected names in an auction.

Claim 11 covers defining a user-selected name specifically to appear like a nested subdomain (e.g., "my.name.university.org").

III. Description of Prior Art Needed to Bust the Patent

To qualify, prior art must have been publicly available before **August 6, 1999**. In order to invalidate the entire patent, we will need prior art that covers each claim. Invalidating claims 1 and 12 is most important, however, because they are the broadest and because some of the dependent claims cover techniques that were well-known in 1999 and are likely to be considered obvious to use in combination with claims 1 and 12. Therefore, we will focus on descriptions of art for those two claims.

As described above, the patent contains both method and system claims. A method claim covers a way of doing something, like a recipe, and a system claim covers a tool for doing it, like a product. For this patent busting project, prior art that shows either will help us defeat the claims.

Any publication (an article, software code, a web page, a patent, a conference poster presentation, or other public writing) or a product can serve as prior art. The prior art should cover all aspects of a claim to defeat that claim. Prior art for a dependent claim must include all aspects of the independent claim (claim 1 or 12) that it's based on, as well as whatever the dependent claim adds.

Ideal prior art for busting claims 1 and 12 would consist of a description of all of the following:

a) a web server/hosting system that uses a wildcard DNS entry *in combination with*:
b) virtual hosting that provides
c) content served from multiple subdomains, where
d) the content is stored using personalized directory names similar to those used in associated subdomains. (e.g., alice.unversity.edu maps to http://www.university.edu/~alice).

The DNS wildcard is essential to the patent, since the inventor was forced to narrow the claim by including the DNS wildcard element in response to an objection from the Patent Office. Use of virtual hosting in which the subdomain name is mapped to a directory with the same name is also essential.[1]

Very strong prior art would be an example of a consumer web-hosting set-up that supported personalized subdomains or something very similar. Such prior art would probably include, in addition to elements (1)-(4), features that allowed the

[1] A virtual hosting set-up in which the directory did not have the same name as the personal subdomain, but was mapped using a database lookup, might be enough to render the invention obvious, so the EFF would be interested in prior art of this kind as well.

user who chose the personal subdomain to upload content (home pages, etc.), and would have a registration phase in which the user chose a personal subdomain name that wasn't already taken. This would invalidate claims 8-9 and 18-19.

Prior art that showed elements (1)-(4) above, as well as HTML frames with advertisements would defeat claims 3-5 and 14-15. Similarly, prior art that showed elements (1)-(4) and also used load-balancing equipment would defeat claims 6-7 and 16-17. Prior art that showed elements (1)-(4) and also showed domains being sold by auction would invalidate claims 10 and 20.

Prior art that showed elements (1)-(4) and also showed subdomains that looked like nested domains would invalidate claim 11. With this claim, the applicant was focusing on subdomains that appeared categorized, like "alice.engineering.university.edu" as a space for people studying engineering. Prior art of this kind would be especially relevant.

Where To Look For Prior Art

Here are a few likely places where prior art systems might have existed:

1) Foreign ISPs
2) Universities with creative computer science students
3) ISPs that provided web hosting services

The reason for the first is simple. Most foreign countries administer their top-level domain as a two-letter country code (e.g., au, de, jp, uk). Then, they duplicate the classic .com/.edu structure. (e.g., edu.au, co.uk). That means that any given company or university was already operating at the third level (e.g., unsw.edu.au, yahoo.co.uk). Thus, any company that was hosting website space on a single server might have been actively using wild-cards in the DNS zone file to direct traffic in a single computer and selling individual directory space. The "user selected label" in the above simply would be the name of the company.

The reason for the second is that savvy CS students (and CS administrators) like cool hacks. They also like to share their knowledge.

The reason for the third is that the alleged inventors described the system using figures that named a web hosting site, WebJump. Naturally, there have been *lots* of web hosting companies. And the more hosting being provided by a company, the more effort that would be required to update DNS files. Thus, the greater the likelihood someone thought of this sooner.

Tip: Check the webpages for your old ISPs on the Wayback Machine at http://web.archive.org. Maybe they offered personal subdomains without charging for a DNS entry.

Where to send information on prior art: priorart at eff dot org or http://www.eff.org/patent/wanted/prior.php?p=test

Further Background

Internet Domains

On the Internet, computers find each other using Internet Protocol (IP) addresses. An IP address is a 32-bit number generally written in the form of four numbers separated by dots, such as 192.0.0.1. These numbers are hard to remember and somewhat static. Thus, a mnemonic translation system was established so that people could use easily remembered, descriptive names to identify computers, such as www.eff.org.

The Domain Name System (DNS) translates these names into an IP address that is used by the underlying IP communications protocol. Because of this translation, it is possible for a single computer to map to numerous domain names, regardless of whether they are similar or not. What matters is that the DNS map have an appropriate entry. For a detailed explanation of DNS, check out http://www.faqs.org/rfcs/rfc1035.html; also http://www.faqs.org/rfcs/rfc1591.html.

Basically, large domains are subdivided into smaller domains that contain one or more computers. You may be familiar with this already, but let's break it down. Some examples of first-level (a/k/a top level) domains are ".com," ".edu," and ".uk." Adding another qualifier in front yields a second-level domain (e.g., second-level.com, university.edu, com.uk). And another would, of course be a third-level domain name (e.g., www.eff.org). It's possible for there to be many levels, but 3-4 is pretty common.

A simplified snippet of a DNS zone file might look like this:

www.university.edu	192.0.0.1
alice.university.edu	192.0.0.2
bob.university.edu	192.0.0.3
charlie.university.edu	192.0.0.4

When a browser processes a URL, it attempts to open a connection to the host computer at the IP address given by the DNS service. Once the connection is open, the browser will send the request URL to the host computer for processing.

Accessing Files over the Internet

"Surfing the web" means getting information stored in a remote computer transferred back to your browser (or other client program) for display. Naturally, these files are stored somewhere (or created on the fly using a script or program), and URLs are a standardized way of requesting a specific file (or resource) from a computer. URLs can be formatted as an index to the file system of a computer, and this allows the system administrator to set up the directories of that computer

in such a way that individual users can control the access to the information stored there. One popular way of locating user directories on a Unix computer is to combine the name of the user with the domain name (using a tilde as a special signal):

http://www.university.edu/~alice/index.html

The above example would open a connection to the computer at 192.0.0.1 and request file "index.html" located in the "alice" web directory. Suppose there were three users with their own websites. In the DNS snippet above, each user (e.g., Alice, Bob, and Charlie) would have a website hosted at three separate IP addresses. However, each user could also be given separate entries in the DNS file that resolve them to a single IP address:

alice.university.edu	192.0.0.1
bob.university.edu	192.0.0.1
charlie.university.edu	192.0.0.1

One downside with this method is that with every new user the DNS file becomes larger and has to be updated. However, another way is to use a special character, such as an asterisk, to act as a wild-card. For example, all requests to specific sub-domains can be allocated to a single computer for processing by using an entry like this:

*.students.university.edu	192.0.0.99
*.faculty.university.edu	192.0.0.98
*.university.edu	192.0.0.1

This alleviates the need for updating the DNS file when new users are added to the system. Also, by using this method, a webserver can provide files from the user directory by mapping the third-level label "alice," for example, back into the traditional tilde form. Thus, the webserver responding to the URL http://alice.university.edu/index.html would serve the same page as the URL http://www.university.edu/~alice/index/html *without* creating separate DNS entries for each individual sub-domain.