total number of requests.)

[0014] Proxy caches depend on cache control hints delivered with resources to determine when the resources should be replaced. These hints are predictive, and are necessarily often incorrect, so proxy caches frequently serve stale data. In many cases, proxy cache operators instruct their proxy to ignore hints in order to make the cache more efficient, even though this causes it to more frequently serve stale data.

[0015] Proxy caches hide the activity of clients from publishers. Once a resource is cached, the publisher has no way of knowing how often it was accessed from the cache.

[0016] WO 97/29423A discloses load balancing across a network of servers according to server load.

Summary of the Invention

[0017] According to the present invention there is provided a system operative in a computer network (100) in which clients (106) connect to an origin server (102), the system comprising: a set of repeater servers (104a...) (also known as content servers), distinct from the origin server (102), for hosting at least some of the embedded objects of web pages that are normally hosted by the origin server (102); a routine for modifying at least one embedded object URL of a web page to resolve to a server in the, set of repeater servers instead of to the origin server; and a server selector mechanism constructed and adapted to identify, for a particular client request, an appropriate repeater server from the set of repeater servers, wherein in response to requests for the web page, generated by the clients, the web page including the modified embedded object URL is served from the origin server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the server selector mechanism; characterised in that the server selector mechanism identifies a repeater server based on the requesting client's location and data relating to network traffic conditions.

[0018] According to a second aspect of the present invention there is provided, in a system operative in a computer network, the system comprising; a set of repeater servers (104a...)(also known as content servers), distinct from an origin server (102), for hosting at least some of the embedded objects of web pages that are normally hosted by the origin server; a routine for causing at least one embedded object URL of a web page to resolve to a server in the set of repeater servers instead of to the origin server; and a server selector mechanism constructed and adapted to identify, for a particular client request, an appropriate repeater server from the set of repeater servers, wherein the server selector mechanism identifies a repeater server based on the requesting client's location and data relating to network traffic conditions, a method of serving a page supported at an origin server, the page comprising a base document (for example, a markup language base document) having associated therewith embedded objects, each embedded object being identified by a uniform resource locator (URL), the method comprising: modifying (preferably re-writing) the URL of an embedded object to generate a modified URL, the modified URL including anew host name referencing a repeater server in the set of repeater servers; in response to a request to serve the page (preferably received at the origin server), serving the page with the modified URL; attempting to serve the embedded object from a repeater server as identified by the new host name; and if a copy of the embedded object is unavailable at the repeater server, obtaining the embedded object from the origin server.

[0019] This invention provides a way for servers in a computer network to off-load their processing of requests for selected resources by determining a different server (a "repeater") to process those requests. The selection of the repeater can be made dynamically, based on information about possible repeaters.

[0020] If a requested resource contains references to other resources, some or all of these references can be replaced by references to repeaters.

[0021] Accordingly, in one aspect, this invention is a method of processing resource requests in a computer network First a client makes a request for a particular resource from an origin server, the request including a resource identifier for the particular resource, the resource identifier sometimes including an indication of the origin server. Requests arriving at the origin server do not always include an indication of the origin server; since they arc sent to the origin server, they do not need to name it. A mechanism referred to as a reflector, co-located with the origin server, intercepts the request from the client to the origin server and decides whether to reflect the request or to handle it locally. If the reflector decides to handle the request locally, it forwards it to the origin server, otherwise it selects a "best" repeater to process the request. If the request is reflected, the client is provided with a modified resource identifier designating the repeater.

[0022] The client gets the modified resource identifier from the reflector and makes a request for the particular resource from the repeater designated in the modified resource identifier.

[0023] When the repeater gets the client's request, it responds by returning the requested resource to the client. If the repeater has a local copy of the resource then it returns that copy, otherwise it forwards the request to the origin server to get the resource, and saves a local copy of the resource in order to serve subsequent requests.

[0024] The selection by the reflector of an appropriate repeater to handle the request can be done in a number of ways. In the preferred embodiment, it is done by first pre-partitioning the network into "cost groups" and then determining which cost group the client is in. Next, from a plurality of repeaters in the network, a set of repeaters

is selected, the members of the set having a low cost relative to the cost group which the client is in. In order to determine the lowest cost, a table is maintained and regularly updated to define the cost between each group and each repeater. Then one member of the set is selected, preferably randomly, as the best repeater.

[0025] If the particular requested resource itself can contain identifiers of other resources, then the resource may be rewritten (before being provided to the client). In particular, the resource is rewritten to replace at least some of the resource identifiers contained therein with modified resource identifiers designating a repeater instead of the origin server. As a consequence of this rewriting process, when the client requests other resources based on identifiers in the particular requested resource, the client will make those requests directly to the selected repeater, bypassing the reflector and origin server entirely.

[0026] Resource rewriting must be performed by reflectors. It may also be performed by repeaters, in the situation where repeaters "peer" with one another and make copies of resources which include rewritten resource identifiers that designate a repeater.

[0027] In a preferred embodiment, the network is the Internet and the resource identifier is a uniform resource locator (URL) for designating resources on the Internet, and the modified resource identifier is a URL designating the repeater and indicating the origin server (as described in step B3 below), and the modified resource identifier is provided to the client using a REDIRECT message. Note, only when the reflector is "reflecting" a request is the modified resource identifier provided using a REDIRECT message.

[0028] In another aspect, this invention is a computer network comprising a plurality of origin servers, at least some of the origin servers having reflectors associated therewith, and a plurality of repeaters.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] The above and other objects and advantages of the invention will be apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, in which the reference characters refer to like parts throughout and in which:

FIGURE 1 depicts a portion of a network environment according to the present invention; and
FIGURES 2-6 are flow charts of the operation of the present invention.

DETAILED DESCRIPTION OF THE PRESENTLY PREFERRED EXEMPLARY EMBODIMENTS

Overview

[0030] **FIGURE 1** shows a portion of a network environment **100** according to the present invention, wherein a mechanism (reflector **108**, described in detail below) at a server (herein origin server **102**) maintains and keeps track of a number of partially replicated servers or repeaters **104a, 104b,** and **104c**. Each repeater **104a, 104b,** and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**. Reflector **108** is connected to a particular repeater known as its "contact" repeater ("Repeater B" **104b** in the system depicted in **FIGURE 1**). Preferably each reflector maintains a connection with a single repeater known as its contact, and each repeater maintains a connection with a special repeater known as its master repeater (e.g., repeater **104m** for repeaters **104a, 104b** and **104c** in **FIGURE 1**).

[0031] Thus, a repeater can be considered as a dedicated proxy server that maintains a partial or sparse mirror of the origin server **102**, by implementing a distributed coherent cache of the origin server. A repeater may maintain a (partial) mirror of more than one origin server. In some embodiments, the network **100** is the Internet and repeaters mirror selected resources provided by origin servers in response to clients' HTTP (hypertext transfer protocol) and FTP (file transfer protocol) requests.

[0032] A client **106** connects, via the network **100**, to origin server **102** and possibly to one or more repeaters **104a** etc.

[0033] Origin server **102** is a server at which resources originate. More generally, the origin server **102** is any process or collection of processes that provide resources in response to requests from a client **106**. Origin server **102** can be any off-the-shelf Web server. In a preferred embodiment, origin server **102** is typically a Web server such as the Apache server or Netscape Communications Corporation's Enterprise™ server.

[0034] Client **106** is a processor requesting resources from origin server **102** on behalf of an end user. The client **106** is typically a user agent (e.g., a Web browser such as Netscape Communications Corporation's Navigator™) or a proxy for a user agent. Components other than the reflector **108** and the repeaters **104a, 104b,** etc., may be implemented using commonly available software programs. In particular, this invention works with any HTTP client (e.g., a Web browser), proxy cache, and Web server. In addition, the reflector **108** might be fully integrated into the data server **112** (for instance, in a Web Server). These components might be loosely integrated based on the use of extension mechanisms (such as so-called add-in modules) or tightly integrated by modifying the service component specifically to support the repeaters.

[0035] Resources originating at the origin server **102** may be static or dynamic. That is, the resources may be fixed or they may be created by the origin server **102** specifically in response to a request. Note that the terms "static" and "dynamic" are relative, since a static resource may change at some regular, albeit long, inter-

val.

[0036] Resource requests from the client **106** to the origin server **102** are intercepted by reflector **108** which for a given request either forwards the request on to the origin server **102** or conditionally reflects it to some repeater **104a, 104b,** etc. in the network **100**. That is, depending on the nature of the request by the client **106** to the origin server **102**, the reflector **108** either serves the request locally (at the origin server **102**), or selects one of the repeaters (preferably the best repeater for the job) and reflects the request to the selected repeater. In other words, the reflector **108** causes requests for resources from origin server **102**, made by client **106**, to be either served locally by the origin server **102** or transparently reflected to the best repeater **104a, 104b,** etc. The notion of a best repeater and the manner in which the best repeater is selected are described in detail below.

[0037] Repeaters **104a, 104b**, etc. are intermediate processors used to service client requests thereby improving performance and reducing costs in the manner described herein. Within repeaters **104a, 104b**, etc., are any processes or collections of processes that deliver resources to the client **106** on behalf of the origin server **102**. A repeater may include a repeater cache **110**, used to avoid unnecessary transactions with the origin server **102**.

[0038] The reflector **108** is a mechanism, preferably a software program, that intercepts requests that would normally be sent directly to the origin server **102**. While shown in the drawings as separate components, the reflector **108** and the origin server **102** are typically co-located, e.g., on a particular system such as data server **112**. (As discussed below, the reflector **108** may even be a "plug in" module that becomes part of the origin server **102**.

[0039] **FIGURE 1** shows only a part of a network **100** according to this invention. A complete operating network consists of any number of clients, repeaters, reflectors, and origin servers. Reflectors communicate with the repeater network, and repeaters in the network communicate with one another.

Uniform Resource Locators

[0040] Each location in a computer network has an address which can generally be specified as a series of names or numbers. In order to access information, an address for that information must be known. For example, on the World Wide Web ("the Web") which is a subset of the Internet, the manner in which information address locations are provided has been standardized into Uniform Resource Locators (URLs). URLs specify the location of resources (information, data files, etc.) on the network.

[0041] The notion of URLs becomes even more useful when hypertext documents are used. A hypertext document is one which includes, within the document itself, links (pointers or references) to the document itself or to other documents. For example, in an on-line legal research system, each case may be presented as a hypertext document. When other cases are cited, links to those cases can be provided. In this way, when a person is reading a case, they can follow cite links to read the appropriate parts of cited cases.

[0042] In the case of the Internet in general and the World Wide Web specifically, documents can be created using a standardized form known as the Hypertext Markup Language (HTML). In HTML, a document consists of data (text, images, sounds, and the like), including links to other sections of the same document or to other documents. The links are generally provided as URLs, and can be in relative or absolute form. Relative URLs simply omit the parts of the URL which are the same as for the document including the link, such as the address of the document (when linking to the same document), etc. In general, a browser program will fill in missing parts of a URL using the corresponding parts from the current document, thereby forming a fully formed URL including a fully qualified domain name, etc.

[0043] A hypertext document may contain any number of links to other documents, and each of those other documents may be on a different server in a different part of the world. For example, a document may contain links to documents in Russia, Africa, China and Australia. A user viewing that document at a particular client can follow any of the links transparently (i.e., without knowing where the document being linked to actually resides). Accordingly, the cost (in terms of time or money or resource allocation) of following one link versus another may be quite significant.

[0044] URLs generally have the following form (defined in detail in T. Berners-Lee et al, *Uniform Resource Locators* (URL), Network Working Group, Request for Comments: 1738, Category: Standards Track, December 1994, located at "http://ds.internic.net/rfc/rfc1738.txt"):

scheme://bost[:port]/url-path

where "scheme" can be a symbol such as *"file"* (for a file on the local system), *"ftp"* (for a file on an anonymous FTP file server), *"http"* (for a file on a file on a Web server), and *"telnet" (for* a connection to a Telnet-based service). Other schemes, can also be used and new schemes are added every now and then. The port number is optional, the system substituting a default port number (depending on the scheme) if none is provided. The "host" field maps to a particular network address for a particular computer. The "url-path" is relative to the computer specified in the "host" field. A url-path is typically, but not necessarily, the pathname of a file in a web server directory.

[0045] For example, the following is a URL identifying a file "F" in the path "A/B/C" on a computer at *"www.uspto.gov"*:

bttp://www.uspto.gov/A/B/C/F

[0046] In order to access the file "F" (the resource)

specified by the above URL, a program (e.g., a browser) running on a user's computer (i.e., a client computer) would have to first locate the computer (i.e., a server computer) specified by the host name. I.e., the program would have to locate the server "*www.uspto.gov*". To do this, it would access a Domain Name Server (DNS), providing the DNS with the host name ("*www.uspto.gov*"). The DNS acts as a kind of centralized directory for resolving addresses from names. If the DNS determines that there is a (remote server) computer corresponding to the name "*www.uspto.gov*", it will provide the program with an actual computer network address for that server computer. On the Internet this is called an Internet Protocol (or IP) address and it has the form "123.345.456.678". The program on the user's (client) computer would then use the actual address to access the remote (server) computer.

[0047] The program opens a connection to the HTTP server.(Web server) on the remote computer "*www.uspto.gov*" and uses the connection to send a request message to the remote computer (using the HTTP scheme). The message is typically an HTTP GET request which includes the url-path of the requested resource, "A/B/C/F". The HTTP server receives the request and uses it to access the resource specified by the url-path "A/B/C/F". The server returns the resource over the same connection.

[0048] Thus, conventionally HTTP client requests for Web resources at an origin server **102** are processed as follows (see **FIGURE 2**) (This is a description of the process when no reflector **108** is installed.):

A1. A browser (e.g., Netscape's Navigator) at the client receives a resource identifier (i.e., a URL) from a user.

A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server. The browser also extracts a port number, if one is present, or uses a default port number (the default port number for http requests is 80).

A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**.

A4. The client **106** then sends a (GET) request over the connection identifying the requested resource.

A5. The origin server **102** receives the request and

A6. locates or composes the corresponding resource.

A7. The origin server **102** then sends back to the client **106** a reply containing the requested resource (or some form of error indicator if the resource is unavailable). The reply is sent to the client over the same connection as that on which the request was received from the client.

A8. The client **106** receives the reply from the origin server **102**.

[0049] There are many variations of this basic model. For example, in one variation, instead of providing the client with the resource, the origin server can tell the client to re-request the resource by another name. To do so, in A7 the server **102** sends back to the client **106** a reply called a "REDIRECT" which contains a new URL indicating the other name. The client **106** then repeats the entire sequence, normally without any user intervention, this time requesting the resource identified by the new URL.

System Operation

[0050] In this invention reflector **108** effectively takes the place of an ordinary Web server or origin server **102**. The reflector **108** does this by taking over the origin server's IP address and port number. In this way, when a client tries to connect to the origin server **102**, it will actually connect to the reflector **108**. The original Web server (or origin server **102**) must then accept requests at a different network (IP) address, or at the same IP address but on a different port number. Thus, using this invention, the server referred to in A3-A7 above is actually a reflector **108**.

[0051] Note that it is also possible to leave the origin server's network address as it is and to let the reflector run at a different address or on a different port. In this way the reflector does not intercept requests sent to the origin server, but can still be sent requests addressed specifically to the reflector. Thus the system can be tested and configured without interrupting its normal operation.

[0052] The reflector **108** supports the processing as follows (see **FIGURE 3**):

upon receipt of a request,

B1 The reflector **108** analyzes the request to determine whether or not to reflect the request. To do this, first the reflector determines whether the sender (client **106**) is a browser or a repeater. Requests issued by repeaters must be served locally by the origin server **102**. This determination can be made by looking up the network (IP) address of the sender in a list of known repeater network (IP) addresses. Alternatively, this determination could be made by attaching information to a request to indicate that the request is from a specific repeater, or repeaters can request re-

sources from a special port other than the one used for ordinary clients.

B2 If the request is not from a repeater, the reflector looks up the requested resource in a table (called the "rule base") to determine whether the resource requested is "repeatable". Based on this determination, the reflector either reflects the request (B3, described below) or serves the request locally (B4, described below).
The rule base is a list of regular expressions and associated attributes. (Regular expressions are well-known in the field of computer science. A small bibliography of their use is found in *Abo, et al.*, "Compilers, Principles, techniques and tools", Addison-Wesley, 1986, pp. 157-158.) The resource identifier (URL) for a given request is looked up in the rule base by matching it sequentially with each regular expression. The first match identifies the attributes for the resource, namely repeatable or local. If there is no match in the rule base, a default attribute is used. Each reflector has its own rule base, which is manually configured by the reflector operator.

B3. To reflect a request, (to serve a request locally go to B4), as shown in **FIGURE 4**, the reflector determines (B3-1) the best repeater to reflect the request to, as described in detail below. The reflector then creates (B3-2) a new resource identifier (URL) (using the requested URL and the best repeater) that identifies the same resource at the selected repeater.
It is necessary that the reflection step create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information. This is done by creating a new URL as follows:

D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format:
http://<repeater>/<server>/<path>
If the form used is other than "http", the preferred embodiment uses the following format:
http://<repeater>/<server>@proxy=<scheme>@/<path>
The reflector can also attach a MIME type to the request, to cause the repeater to provide that MIME type with the result. This is useful because many protocols (such as FTP) do not provide a way to attach a MIME type to a resource. The format is
http://<repeater>/<server>@proxy=<scheme>:<type>@/<path>
This URL is interpreted when received by the repeater.
The reflector then sends (B3-3) a REDIRECT reply containing this new URL to the requesting client. The HTTP REDIRECT command allows the reflector to send the browser a single URL to retry the request.

B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. In this mode, the reflector acts as a reverse proxy server. The origin server **102** processes the request in the normal manner (A5-A7). The reflector then obtains the origin server's reply to the request which it inspects to determine if the requested resource is an HTML document, i.e., whether the requested resource is one which itself contains resource identifiers.

B5. If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it, as described below. The resource, possibly as modified by rewriting, is then returned in a reply to the requesting client **106**.
If the requesting client is a repeater, the reflector may temporarily disable any cache-control modifiers which the origin server attached to the reply. These disabled cache-control modifiers are later re-enabled when the content is served from the repeater. This mechanism makes it possible for the origin server to prevent resources from being cached at normal proxy caches, without affecting the behavior of the cache at the repeater.

B6. Whether the request is reflected or handled locally, details about the transaction, such as the current time, the address of the requester, the URL requested, and the type of response generated, are written by the reflector to a local log file.

[0053] By using a rule base (B2), it is possible to selectively reflect resources. There are a number of reasons that certain particular resources cannot be effectively repeated (and therefore should not be reflected), for instance:

- the resource is composed uniquely for each request;
- the resource relies on a so-called cookie (browsers will not send cookies to repeaters with different domain names);
- the resource is actually a program (such as a Java applet) that will run on the client and that wishes to connect to a service (Java requires that the service be running on the same machine that provided the applet).

[0054] In addition, the reflector **108** can be configured

so that requests from certain network addresses (e.g., requests from clients on the same local area network as the reflector itself) are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate, as described below.

[0055] A request which is reflected is automatically mirrored at the repeater when the repeater receives and processes the request.

[0056] The combination of the reflection process described here and the caching process described below effectively creates a system in which repeatable resources are migrated to and mirrored at the selected reflector, while non-repeatable resources are not mirrored.

Alternate Approach

[0057] Placing the origin server name in the reflected URL is generally a good strategy, but it may be considered undesirable for aesthetic or (in the case, e.g., of cookies) certain technical reasons.

[0058] It is possible to avoid the need for placing both the repeater name and the server name in the URL. Instead, a "family" of names may be created for a given origin server, each name identifying one of the repeaters used by that server.

[0059] For instance, if www.example.com is the origin server, names for three repeaters might be created:

wr1.example.com
wr2.example.com
wr3.example.com

[0060] The name "wr1.example.com" would be an alias for repeater 1, which might also be known by other names such as "wr1.anotherExample.com" and "wr1.example.edu".

[0061] If the repeater can determine by which name it was addressed, it can use this information (along with a table that associates repeater alias names with origin server names) to determine which origin server is being addressed. For instance, if repeater 1 is addressed as wr1.example.com, then the origin server is "www.example.com"; if it is addressed as "wr1.anotherExample.com", then the origin server is "www.anotherExample.com".

[0062] The repeater can use two mechanisms to determine by which alias it is addressed:

1. Each alias can be associated with a different IP address. Unfortunately, this solution does not scale well, as IP addresses are currently scarce, and the number of IP addresses required grows as the product of origin servers and repeaters.
2. The repeater can attempt to determine the alias name used by inspecting the "host:" tag in the HTTP header of the request. Unfortunately, some old browsers still in use do not attach the "host:" tag to a request. Reflectors would need to identify such browsers (the browser identity is a part of each request) and avoid this form of reflection.

How a Repeater Handles a Request

[0063] When a browser receives a REDIRECT response (as produced in B3), it reissues a request for the resource using the new resource identifier (URL) (A1-A5). Because the new identifier refers to a repeater instead of the origin server, the browser now sends a request for the resource to the repeater which processes a request as follows, with reference to **FIGURE 5**:

C1. First the repeater analyzes the request to determine the network address of the requesting client and the path of the resource requested. Included in the path is an origin server name (as described above with reference to B3).

C2. The repeater uses an internal table to verify that the origin server belongs to a known "subscriber". A subscriber is an entity (e.g., a company) that publishes resources (e.g., files) via one or more origin servers. When the entity subscribes, it is permitted to utilize the repeater network. The subscriber tables described below include the information that is used to link reflectors to subscribers.
If the request is not for a resource from a known subscriber, the request is rejected. To reject a request, the repeater returns a reply indicating that the requested resource does not exist.

C3. The repeater then determines whether the requested resource is cached locally. If the requested resource is in the repeater's cache it is retrieved. On the other hand, if a valid copy of the requested resource is not in the repeater's cache, the repeater modifies the incoming URL, creating a request that it issues directly to the originating reflector which processes it (as in B1-B6). Because this request to the originating reflector is from a repeater, the reflector always returns the requested resource rather than reflecting the request. (Recall that reflectors always handle requests from repeaters locally.) If the repeater obtained the resource from the origin server, the repeater then caches the resource locally.
If a resource is not cached locally, the cache can query its "peer caches" to see if one of them contains the resource, before or at the same time as requesting the resource from the reflector/origin server. If a peer cache responds positively in a limited period of time (preferably a small fraction of a second), the resource will be retrieved from

the peer cache.

C4. The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client.

C5. Details about the transaction, such as the associated reflector, the current time, the address of the requester, the URL requested, and the type of response generated, are written to a local log file at the repeater.

[0064] Note that the bottom row of **FIGURE 2** refers to an origin server, or a reflector, or a repeater, depending on what the URL in step A1 identifies.

Selecting the Best Repeater

[0065] If the reflector **108** determines that it will reflect the request, it must then select the best repeater to handle that request (as referred to in step B3-1). This selection is performed by the Best Repeater Selector (BRS) mechanism described here.

[0066] The goal of the BRS is to select, quickly and heuristically, an appropriate repeater for a given client given only the network address of the client. An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance. The mechanism used here relies on specific, compact, pre-computed data to make a fast decision. Other, dynamic solutions can also be used to select an appropriate repeater.

[0067] The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. These three tables (described below) are computed off-line and downloaded to each reflector by its contact in the repeater network.

[0068] **The Group Reduction Table** places every network address into a group, with the goal that addresses in a group share relative costs, so that they would have the same best repeater under varying conditions (i.e., the BRS is invariant over the members of the group).

[0069] **The Link Cost Table** is a two dimensional matrix which specifies the current cost between each repeater and each group. Initially, the link cost between a repeater and a group is defined as the "normalized link cost" between the repeater and the group, as defined below. Over time, the table will be updated with measurements which more accurately reflect the relative cost of transmitting a file between the repeater and a member of the group. The format of the Link Cost Table is <Group ID> <Group ID> <link cost>, where the Group ID's are given as AS numbers.

[0070] **The Load Table** is a one dimensional table which identifies the current load at each repeater. Because repeaters may have different capacities, the load is a value that represents the ability of a given repeater to accept additional work. Each repeater sends its current load to a central master repeater at regular intervals, preferably at least approximately once a minute. The master repeater broadcasts the load Table to each reflector in the network, via the contact repeater.

[0071] A reflector is provided entries in the Load Table only for repeaters which it is assigned to use. The assignment of repeaters to reflectors is performed centrally by a repeater network operator at the master repeater. This assignment makes it possible to modify the service level of a given reflector. For instance, a very active reflector may use many repeaters, whereas a relatively inactive reflector may use few repeaters.

[0072] Tables may also be configured to provide selective repeater service to subscribers in other ways, e. g., for their clients in specific geographic regions, such as Europe or Asia.

Measuring Load

[0073] In the presently preferred embodiments, repeater load is measured in two dimensions, namely

1. requests received by the repeater per time interval *(RRPT)*, and
2. bytes sent by the repeater per time interval *(BSPT)*.

[0074] For each of these dimensions, a maximum capacity setting is set. The maximum capacity indicates the point at which the repeater is considered to be fully loaded. A higher RRPT capacity generally indicates a faster processor, whereas a higher BSPT capacity generally indicates a wider network pipe. This form of load measurement assumes that a given server is dedicated to the task of repeating.

[0075] Each repeater regularly calculates its current RRPT and BSPT, by accumulating the number of requests received and bytes sent over a short time interval. These measurements are used to determine the repeater's load in each of these dimensions. If a repeater's load exceeds its configured capacity, an alarm message is sent to the repeater network administrator.

[0076] The two current load components are combined into a single value indicating overall current load. Similarly, the two maximum capacity components are combined into a single value indicating overall maximum capacity. The components are combined as follows:

$$\text{current-load} = B \times \text{current RRPT} + (1 - B) \times \text{current BSPT}$$

$$\text{max-load} = B \times \text{max RRPT} + (1 - B) \times \text{max BSPT}$$

[0077] The factor *B*, a value between 0 and 1, allows the relative weights of RRPT and BSPT to be adjusted, which favors consideration of either processing power or bandwidth.

[0078] The overall current load and overall maximum capacity values are periodically sent from each repeater to the master repeater, where they are aggregated in the Load Table, a table summarizing the overall load for all repeaters. Changes in the Load Table are distributed automatically to each reflector.

[0079] While the preferred embodiment uses a two-dimensional measure of repeater load, any other measure of load can be used.

Combining Link Costs and Load

[0080] The BRS computes the cost of servicing a given client from each eligible repeater. The cost is computed by combining the available capacity of the candidate repeater with the cost of the link between that repeater and the client. The link cost is computed by simply looking it up in the link Cost table.

[0081] The cost is determined using the following formula:

$$threshold = K * max\text{-}load$$

$$capacity = max(max\text{-}load - current\text{-}load, e)$$

$$capacity = min(capacity, threshold)$$

$$cost = link\text{-}cost * threshold / capacity$$

[0082] In this formula, *e* is a very small number (epsilon) and *K* is a tuning factor initial set to 0.5. This formula causes the cost to a given repeater to be increased, at a rate defined by *K*, if its capacity falls below a configurable threshold.

[0083] Given the cost of each candidate repeater, the BRS selects all repeaters within a delta factor of the best score. From this set, the result is selected at random.

[0084] The delta factor prevents the BRS from repeatedly selecting a single repeater when scores are similar. It is generally required because available information about load and link costs loses accuracy over time. This factor is tunable.

Best Repeater Selector (BRS)

[0085] The BRS operates as follows, with reference to **FIGURE 6**:

Given a client network address and the three tables described above:

E1. Determine which group the client is in using the Group Reduction Table.

E2. For each repeater in the Link Cost Table and Load Table, determine that repeater's combined cost as follows:

- E2a. Determine the maximum and current load on the repeater (using the Load Table).
- E2b. Determine the link cost between the repeater and the client's group (using the Link Cost Table).
- E2c. Determine the combined cost as described above.

E3. Select a small set of repeaters with the lowest cost.

E4. Select a random member from the set.

[0086] Preferably the results of the BRS processing are maintained in a local cache at the reflector **108**. Thus, if the best repeater has recently been determined for a given client (i.e., for a given network address), that best repeater can be reused quickly without being re-determined. Since the calculation described above is based on statically, pre-computed tables, if the tables have not changed then there is no need to re-determine the best repeater.

Determining the Group Reduction and Link Cost Tables

[0087] The Group Reduction Table and Link Cost Table used in BRS processing are created and regularly updated by an independent procedure referred to herein as *NetMap*. The *NetMap* procedure is run by executing several phases (described below) as needed.

[0088] The term *Group* is used here to refers to an IP "address group".

[0089] The term *Repeater Group* refers to a Group that contains the IP address of a repeater.

[0090] The term *link cost* refers to a statically determined cost for transmitting data between two Groups. In a presently preferred implementation, this is the minimum of the sums of the costs of the links along each path between them. The link costs of primary concern here are link costs between a Group and a Repeater Group.

[0091] The term *relative link cost* refers to the link cost relative to other link costs for the same Group which is calculated by subtracting the minimum link cost from a Group to any Repeater Group from each of its link costs to a Repeater Group.

The term Cost Set refers to a set of Groups that are equivalent in regard to Best Repeater Selection. That is, given the information available, the same repeater

would be selected for any of them.

[0092] The *NetMap* procedure first processes input files to create an internal database called the Group Registry. These input files describe groups, the IP addresses within groups, and links between groups, and come a variety of sources, including publicly available Internet Routing Registry (IRR) databases, BGP router tables, and probe services that are located at various points around the Internet and use publicly available tools (such as "traceroute") to sample data paths. Once this processing is complete, the Group Registry contains essential information used for further processing, namely (1) the identity of each group, (2) the set of IP addresses in a given group, (3) the presence of links between groups indicating paths over which information may travel, and (4) the cost of sending data over a given link.

[0093] The following processes are then performed on the Group Registry file.

Calculate Repeater Group link costs

[0094] The *NetMap* procedure calculates a "link cost" for transmission of data between each Repeater Group and each Group in the Group Registry. This overall link cost is defined as the minimum cost of any path between the two groups, where the cost of a path is equal to the sum of the costs of the individual links in the path. The link cost algorithm presented below is essentially the same as algorithm #562 from ACM journal Transactions on Mathematical Software: "Shortest Path From a Specific Node to All Other Nodes in a Network" by U. Pape, ACM TOMS 6 (1980) pp. 450-455, http://www.netlib.org/toms/562.

[0095] In this processing, the term Repeater Group refers to a Group that contains the IP address of a repeater. A group is a neighbor of another group if the Group Registry indicates that there is a link between the two groups.

[0096] For each target Repeater Group **T**:

- Initialize the link cost between **T** and itself to zero.
- Initialize the link cost between **T** and every other Group to infinity.
- Create a list **L** that will contain Groups that are equidistant from the target Repeater Group **T**.
- Initialize the list **L** to contain just the target Repeater Group **T** itself.
- While the list **L** is not empty:
    - Create an empty list **L'** of neighbors of members of the list **L**.
    - For each Group **G** in the list **L**:
        - For each Group **N** that is a neighbor of **G**:
            - Let cost refer to the sum of the link cost between **T** and **G**, and the link cost between **G** and **N**. The cost between **T** and **G** was determined in the previous pass of the algorithm; the link cost between **G** and **N** is from the Group Registry.
            - If cost is less than the link cost between **T** and **N**:
                - Set the link cost between **T** and **N** to cost.
                - Add **N** to **L'** if it is not already on it.
    - Set **L** to **L'**.

Calculate Cost Sets

[0097] A Cost Set is a set of Groups that are equivalent with respect to Best Repeater Selection. That is, given the information available, the same repeater would be selected for any of them.

[0098] The "cost profile" of a Group G is defined herein as the set of costs between G and each Repeater. Two cost profiles are said to be equivalent if the values in one profile differ from the corresponding values in the other profile by a constant amount.

[0099] Once a client Group is known, the Best Repeater Selection algorithm relies on the cost profile for information about the Group. If two cost profiles are equivalent, the BRS algorithm would select the same repeater given either profile.

[0100] A Cost Set is then a set of groups that have equivalent cost profiles.

[0101] The effectiveness of this method can be seen, for example, in the case where all paths to a Repeater from some Group A pass through some other Group B. The two Groups have equivalent cost profiles (and are therefore in the same Cost Set) since whatever Repeater is best for Group A is also going to be best for Group B, regardless of what path is taken between the two Groups.

[0102] By normalizing cost profiles, equivalent cost profiles can be made identical. A normalized cost profile is a cost profile in which the minimum cost has the value zero. A normalized cost profile is computed by finding the minimum cost in the profile, and subtracting that value from each cost in the profile.

[0103] Cost Sets are then computed using the following algorithm:

- For each Group **G**:
    - Calculate the normalized cost profile for **G**
    - Look for a Cost Set with the same normalized cost profile.
    - If such as set is found, add **G** to the existing Cost Set;
    - otherwise, create a new Cost Set with the calculated normalized cost profile, containing only

G.

[0104] The algorithm for finding Cost Sets employs a hash table to reduce the time necessary to determine whether the desired Cost Set already exists. The hash table uses a hash value computed from cost profile of G.

[0105] Each Cost Set is then numbered with a unique Cost Sent Index number. Cost Sets are then used in a straightforward manner to generate the Link Cost Table, which gives the cost from each Cost Set to each Repeater.

[0106] As described below, the Group Reduction Table maps every IP address to one of these Cost Sets.

Build IP Map

[0107] The IP Map is a sorted list of records which map IP address ranges to Link Cost Table keys. The format of the IP map is:

<base IP address> <max IP address> <Link Cost Table key>

where IP addresses are presently represented by 32-bit integers. The entries are sorted by descending base address, and by ascending maximum address among equal base addresses, and by ascending link Cost Table key among equal base addresses and maximum addresses. Note that ranges may overlap.

[0108] The *NetMap* procedure generates an intermediate IP map containing a map between IP address ranges and Cost Set numbers as follows:

- For each Cost Set **S**:
    - For each Group **G** in **S**:
        - For each IP address range in **G**:
            - Add a triple (low address, high address, Cost Set number of S) to the IP map.

[0109] The IP map file is then sorted by descending base address, and by ascending maximum address among equal base addresses, and by ascending Cost Set number among equal base addresses and maximum addresses. The sort order for the base address and maximum address minimizes the time to build the Group Reduction Table and produces the proper results for overlapping entries.

[0110] Finally, the *NetMap* procedure creates the Group Reduction Table by processing the sorted IP map. The Group Reduction Table maps IP addresses (specified by ranges) into Cost Set numbers. Special processing of the IP map file is required in order to detect overlapping address ranges, and to merge adjacent address ranges in order to minimize the size of the Group Reduction Table.

[0111] An ordered list of address range segments is maintained, each segment consisting of a base address **B** and a Cost Set number **N**, sorted by base address **B**. (The maximum address of a segment is the base address of the next segment minus one.)

[0112] The following algorithm is used:

- Initialize the list with the elements [-infinity, NOGROUP], [+infinity, NOGROUP].
    - For each entry in the IP map, in sorted order, consisting of (b, m, s),
        - Insert (b, m, s) in the list (recall that IP map entries are of the form (low address, high address Cost Set number of S))
    - For each reserved LAN address range (b, m):

      Insert (b, m, LOCAL) in the list.
    - For each Repeater at address a:

      Insert (a, a, REPEATER) in the list.
    - For each segment **S** in the ordered list:
        - Merge S with following segments with the same Cost Set
        - Create a Group Reduction Table entry with base address from the base address of **S**,
            - max address = next segment's base - 1,
            - group ID = Cost Set number of **S**.

[0113] A reserved LAN address range is an address range reserved for use by LANs which should not appear as a global Internet address. LOCAL is a special Cost Set index different from all others, indicating that the range maps to a client which should never be reflected. REPEATER is a special Cost Set index different from all others, indicating that the address range maps to a repeater. NOGROUP is a special Cost Set index different from all others, indicating that this range of addresses has no known mapping.

[0114] Given (B, M, N), insert an entry in the ordered address list as follows:

Find the last segment (AB, AN) for which AB is less than or equal to B.
If AB is less than B, insert a new segment (B, N) after (AB, AN).
Find the last segment (YB, YN) for which YB is less than or equal to M.
Replace by (XB, N) any segment (XB, NOGROUP) for which XB is greater than B and less than YB.
If YN is not N, and either YN is NOGROUP or YB is less than or equal to B,

Let (ZB, ZN) be the segment following (YB, YN).
If M+1 is less than ZB, insert a new segment (M+1, YN) before (ZB, ZN).
Replace (YB, YN) by (YB, N).

Rewriting HTML Resources

[0115] As explained above with reference to **FIGURE 3** (B5), when a reflector or repeater serves a resource which itself includes resource identifiers (e.g., a HTML resource), that resource is modified (rewritten) to pre-reflect resource identifiers (URLs) of repeatable resources that appear in the resource. Rewriting ensures that when a browser requests repeatable resources identified by the requested resource, it gets them from a repeater without going back to the origin server, but when it requests non-repeatable resources identified by the requested resource, it will go directly to the origin server. Without this optimization, the browser would either make all requests at the origin server (increasing traffic at the origin server and necessitating far more redirections from the origin server), or it would make all requests at the repeater (causing the repeater to redundantly request and copy resources which could not be cached, increasing the overhead of serving such resources).

[0116] Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector). Rewriting uses a so-called BASE directive. The BASE directive lets the HTML identify a different base server. (The base address is normally the address of the HTML resource.)

[0117] Rewriting is performed as follows:

F1. A BASE directive is added at the beginning of the HTML resource, or modified where necessary. Normally, a browser interprets relative URLs as being relative to the default base address, namely, the URL of the HTML resource (page) in which they are encountered. The BASE address added specifies the resource at the reflector which originally served the resource. This means that unprocessed relative URLs (such as those generated by Javascript™ programs) will be interpreted as relative to the reflector. Without this BASE address, browsers would combine relative addresses with repeater names to create URLs which were not in the form required by repeaters (as described above in step D1).

F2. The rewriter identifies directives, such as embedded images and anchors, containing URLs. If the rewriter is running in a reflector, it must parse the HTML file to identify these directives.
If it is running in a repeater, the rewriter may have access to pre-computed information that identifies the location of each URL (placed in the HTML file in step F4).

F3. For each URL encountered in the resource to be re-written, the rewriter must determine whether the URL is repeatable (as in steps B1-B2). If the URL is not repeatable, it is not modified. On the other hand, if the URL is repeatable, it is modified to refer to the selected repeater.

F4. After all URLs have been identified and modified, if the resource is being served to a repeater, a table is appended at the beginning of the resource that identifies the location and content of each URL encountered in the resource. (This step is an optimization which eliminates the need for parsing HTML resources at the repeater.)

F5. Once all changes have been identified, a new length is computed for the resource (page). The length is inserted in the HTTP header prior to serving the resource.

[0118] An extension of HTML, known as XML, is currently being developed. The process of rewriting URLs will be similar for XML, with some differences in the mechanism that parses the resource and identifies embedded URLs.

Handling Non-HTTP Protocols

[0119] This invention makes it possible to reflect references to resources that are served by protocols other than HTTP, for instance, the File Transfer Protocol (FTP) and audio/video stream protocols. However, many protocols do not provide the ability to redirect requests. It is, however, possible to redirect references before requests are actually made by rewriting URLs embedded in HTML pages. The following modifications to the above algorithms are used to support this capability.

[0120] In F4, the rewriter rewrites URLs for servers if those servers appear in a configurable table of cooperating origin server or so-called co-servers. The reflector operator can define this table to include FTP servers and other servers. A rewritten URL that refers to a non-HTTP resource takes the form:

http://<repeater>/<origin server>@proxy=<scheme>[:<type>]@/resource

where <scheme> is a supported protocol name such as "ftp". This URL format is an alternative to the form shown in B3.

[0121] In C3, the repeater looks for a protocol embedded in the arriving request. If a protocol is present and the requested resource is not already cached, the repeater uses the selected protocol instead of the default HTTP protocol to request the resource when serving it and storing it in the cache.

System Configuration and Management

[0122] In addition to the processing described above, the repeater network requires various mechanisms for system configuration and network management. Some of these mechanisms are described here.

[0123] Reflectors allow their operators to synchronize repeater caches by performing publishing operations. The process of keeping repeater caches synchronized is described below. Publishing indicates that a resource or collection of resources has changed.

[0124] Repeaters and reflectors participate in various types of log processing. The results of logs collected at repeaters are collected and merged with logs collected at reflectors, as described below.

Adding Subscribers to the Repeater Network

[0125] When a new subscriber is added to the network, information about the subscriber is entered in a Subscriber Table at the master repeater and propagated to all repeaters in the network. This information includes the *Committed Aggregate Information Rate* (CAIR) for servers belonging to the subscriber, and a list of the repeaters that may be used by servers belonging to the subscriber.

Adding Reflectors to the Repeater Network

[0126] When a new reflector is added to the network, it simply connects to and announces itself to a contact repeater, preferably using a securely encrypted certificate including the repeater's subscriber identifier.

[0127] The contact repeater determines whether the reflector network address is permitted for this subscriber. If it is, the contact repeater accepts the connection and updates the reflector with all necessary tables (using version numbers to determine which tables are out of date).

[0128] The reflector processes requests during this time, but is not "enabled" (allowed to reflect requests) until all of its tables are current.

Keeping Repeater Caches Synchronized

[0129] Repeater caches are coherent, in the sense that when a change to a resource is identified by a reflector, all repeater caches are notified, and accept the change in a single transaction.

[0130] Only the identifier of the changed resource (and not the entire resource) is transmitted to the repeaters; the identifier is used to effectively invalidate the corresponding cached resource at the repeater. This process is far more efficient than broadcasting the content of the changed resource to each repeater.

[0131] A repeater will load the newly modified resource the next time it is requested.

[0132] A resource change is identified at the reflector either manually by the operator, or through a script when files are installed on the server, or automatically through a change detection mechanism (e.g., a separate process that checks regularly for changes).

[0133] A resource change causes the reflector to send an "invalidate" message to its contact repeater, which forwards the message to the master repeater. The invalidate message contains a list of resource identifiers (or regular expressions identifying patterns of resource identifiers) that have changed. (Regular expressions are used to invalidate a directory or an entire server.) The repeater network uses a two-phase commit process to ensure that all repeaters correctly invalidate a given resource.

[0134] The invalidation process operates as follows:

[0135] The master broadcasts a "phase 1" invalidation request to all repeaters indicating the resources and regular expressions describing sets of resources to be invalidated.

[0136] When each repeater receives the phase 1 message, it first places the resource identifiers or regular expressions into a list of resource identifiers pending invalidation.

[0137] Any resource requested (in C3) that is in the pending invalidation list may not be served from the cache. This prevents the cache from requesting the resource from a peer cache which may not have received an invalidation notice. Were it to request a resource in this manner, it might replace the newly invalidated resource by the same, now stale, data.

[0138] The repeater then compares the resource identifier of each resource in its cache against the resource identifiers and regular expressions in the list.

[0139] Each match is invalidated by marking it stale and optionally removing it from the cache. This means that a future request for the resource will cause it to retrieve a new copy of the resource from the reflector.

[0140] When the repeater has completed the invalidation, it returns an acknowledgment to the master. The master waits until all repeaters have acknowledged the invalidation request.

[0141] If a repeater fails to acknowledge within a given period, it is disconnected from the master repeater. When it reconnects, it will be told to flush its entire cache, which will eliminate any consistency problem. (To avoid flushing the entire cache, the master could keep a log of all invalidations performed, sorted by date, and flush only files invalidated since the last time the reconnecting repeater successfully completed an invalidation. In the presently preferred embodiments this is not done since it is believed that repeaters will seldom disconnect.)

[0142] When all repeaters have acknowledged invalidation (or timed out) the repeater broadcasts a "phase 2" invalidation request to all repeaters. This causes the repeaters to remove the corresponding resource identifiers and regular expressions from the list of resource identifiers pending invalidation.

[0143] In another embodiment, the invalidation re-

quest will be extended to allow a "server push". In such requests, after phase 2 of the invalidation process has completed, the repeater receiving the invalidation request will immediately request a new copy of the invalidated resource to place in its cache.

Logs and Log Processing

[0144] Web server activity logs are fundamental to monitoring the activity in a Web site. This invention creates "merged logs" that combine the activity at reflectors with the activity at repeaters, so that a single activity log appears at the origin server showing all Web resource requests made on behalf of that site at any repeater.

[0145] This merged log can be processed by standard processing tools, as if it had been generated locally.

[0146] On a periodic basis, the master repeater (or its delegate) collects logs from each repeater. The logs collected are merged, sorted by reflector identifier and timestamp, and stored in a dated file on a per-reflector basis. The merged log for a given reflector represents the activity of all repeaters on behalf of that reflector. On a periodic basis, as configured by the reflector operator, a reflector contacts the master repeater to request its merged logs. It downloads these and merges them with its locally maintained logs, sorting by timestamp. The result is a merged log that represents all activity on behalf of repeaters and the given reflector.

[0147] Activity logs are optionally extended with information important to the repeater network, if the reflector is configured to do so by the reflector operator. In particular, an "extended status code" indicates information about each request, such as:

1. request was served by a reflector locally;
2. request was reflected to a repeater;*
3. request was served by a reflector to a repeater;*
4. request for non-repeatable resource was served by repeater;*
5. request was served by a repeater from the cache;
6. request was served by a repeater after filling cache;
7. request pending invalidation was served by a repeater.

(The activities marked with "*" represent intermediate states of a request and do not normally appear in a final activity log.)

[0148] In addition, activity logs contain a duration, and extended precision timestamps. The duration makes it possible to analyze the time required to serve a resource, the bandwidth used, the number of requests handled in parallel at a given time, and other quite useful information. The extended precision timestamp makes it possible to accurately merge activity logs.

[0149] Repeaters use the Network Time Protocol (NTP) to maintain synchronized clocks. Reflectors may either use NTP or calculate a time bias to provide roughly accurate timestamps relative to their contact repeater.

Enforcing Committed Aggregate Information Rate

[0150] The repeater network monitors and limits the aggregate rate at which data is served on behalf of a given subscriber by all repeaters. This mechanism provides the following benefits:

1. provides a means of pricing repeater service;
2. provides a means for estimating and reserving capacity, at repeaters;
3. provides a means for preventing clients of a busy site from limiting access to other sites.

[0151] For each subscriber, a "threshold aggregate information rate" (TAIR) is configured and maintained at the master repeater. This threshold is not necessarily the committed rate, it may be a multiple of committed rate, based on a pricing policy.

[0152] Each repeater measures the information rate component of each reflector for which it serves resources, periodically (typically about once a minute), by recording the number of bytes transmitted on behalf of that reflector each time a request is delivered. The table thus created is sent to the master repeater once per period. The master repeater combines the tables from each repeater, summing the measured information of each reflector over all repeaters that serve resources for that reflector, to determine the "measured aggregate information rate" (MAIR) for each reflector.

[0153] If the MAIR for a given reflector is greater than the TAIR for that reflector, the MAIR is transmitted by the master to all repeaters and to the respective reflector.

[0154] When a reflector receives a request, it determines whether its most recently calculated MAIR is greater than its TAIR. If this is the case, the reflector probabilistically decides whether to suppress reflection, by serving the request locally (in B2). The probability of suppressing the reflection increases as an exponential function of the difference between the MAIR and the CAIR.

[0155] Serving a request locally during a peak period may strain the local origin server, but it prevents this subscriber from taking more than allocated bandwidth from the shared repeater network.

[0156] When a repeater receives a request for a given subscriber (in C2), it determines whether the subscriber is running near its threshold aggregate information rate. If this is the case, it probabilistically decides whether to reduce its load by redirecting the request back to the reflector. The probability increases exponentially as the reflector's aggregate information rate approaches its limit.

[0157] If a request is reflected back to a reflector, a special character string is attached to the resource identifier so that the receiving reflector will not attempt to re-

flect it again. In the current system, this string has the form

"src=overload".

[0158] The reflector tests for this string in B2.

[0159] The mechanism for limiting Aggregate Information Rate described above is fairly coarse. It limits at the level of sessions with clients (since once a client has been reflected to a given repeater, the rewriting process tends to keep the client coming back to that repeater) and, at best, individual requests for resources. A more fine-grained mechanism for enforcing TAIR limits within repeaters operates by reducing the bandwidth consumption of a busy subscriber when other subscribers are competing for bandwidth.

[0160] The fine-grained mechanism is a form of data "rate shaping". It extends the mechanism that copies resource data to a connection when a reply is being sent to a client. When an output channel is established at the time a request is received, the repeater identifies which subscriber the channel is operating for, in C2, and records the subscriber in a data field associated with the channel. Each time a "write" operation is about to be made to the channel, the Metered Output Stream first inspects the current values of the MAIR and TAIR, calculated above, for the given subscriber. If the MAIR is larger than the TAIR, then the mechanism pauses briefly before performing the write operation. The length of the pause is proportional to the amount the MAIR exceeds the TAIR. The pause ensures that tasks sending other resources to other clients, perhaps on behalf of other subscribers, will have an opportunity to send their data.

Repeater Network Resilience

[0161] The repeater network is capable of recovering when a repeater or network connection fails.

[0162] A repeater cannot operate unless it is connected to the master repeater. The master repeater exchanges critical information with other repeaters, including information about repeater load, aggregate information rate, subscribers, and link cost.

[0163] If a master fails, a "succession" process ensures that another repeater will take over the role of master, and the network as a whole will remain operational. If a master fails, or a connection to a master fails through a network problem, any repeater attempting to communicate with the master will detect the failure, either through an indication from TCP/IP, or by timing out from a regular "heartbeat" message it sends to the master.

[0164] When any repeater is disconnected from its master, it immediately tries to reconnect to a series of potential masters based on a configurable file called its "succession list".

[0165] The repeater tries each system on the list in succession until it successfully connects to a master. If in this process; it comes to its own name, it takes on the role of master, and accepts connections from other repeaters. If a repeater which is not at the top of the list becomes the master, it is called the "temporary master".

[0166] A network partition may cause two groups of repeaters each to elect a master. When the partition is corrected, it is necessary that the more senior master take over the network. Therefore, when a repeater is temporary master, it regularly tries to reconnect to any master above it in the succession list. If it succeeds, it immediately disconnects from all of the repeaters connected to it. When they retry their succession lists, they will connect to the more senior master repeater.

[0167] To prevent losses of data, a temporary master does not accept configuration changes and does not process log files. In order to take on these tasks, it must be informed that it is primary master by manual modification of its successor list. Each repeater regularly reloads its successor list to determine whether it should change its idea of who the master is.

[0168] If a repeater is disconnected from the master, it must resynchronize its cache when it reconnects to the master. The master can maintain a list of recent cache invalidations and send to the repeater any invalidations it was not able to process while disconnected. If this list is not available for some reason (for instance, because the reflector has been disconnected too long), the reflector must invalidate its entire cache.

[0169] A reflector is not permitted to reflect requests unless it is connected to a repeater. The reflector relies on its contact repeater for critical information, such as load and Link Cost Tables, and current aggregate information rate. A reflector that is not connected to a repeater can continue to receive requests and handle them locally.

[0170] If a reflector loses its connection with a repeater, due to a repeater failure or network outage, it continues to operate while it tries to connect to a repeater.

[0171] Each time a reflector attempts to connect to a repeater, it uses DNS to identify a set of candidate repeaters given a domain name that represents the repeater network. The reflector tries each repeater in this set until it makes a successful contact. Until a successful contact is made, the reflector serves all requests locally. When a reflector connects to a repeater, the repeater can tell it to attempt to contact a different repeater; this allows the repeater network to ensure that no individual repeater has too many contacts.

[0172] When contact is made, the reflector provides the version number of each of its tables to its contact repeater. The repeater then decides which tables should be updated and sends appropriate updates to the reflector. Once all tables have been updated, the repeater notifies the reflector that it may now start reflecting requests.

Using a Proxy Cache within a Repeater

[0173] Repeaters are intentionally designed so that any proxy cache can be used as a component within

them. This is possible because the repeater receives HTTP requests and converts them to a form recognized by the proxy cache.

[0174] On the other hand, several modifications to a standard proxy cache have been or may be made as optimizations. This includes, in particular, the ability to conveniently invalidate a resource, the ability to support cache quotas, and the ability to avoid making an extra copy of each resource as it passes from the proxy cache through the repeater to the requester.

[0175] In a preferred embodiment, a proxy cache is used to implement C3. The proxy cache is dedicated for use only by one or more repeaters. Each repeater requiring a resource from the proxy cache constructs a proxy request from the inbound resource request A normal HTTP GET request to a server contains only the pathname part of the URL-the scheme and server name are implicit. (In an HTTP GET request to a repeater, the pathname part of the URL includes the name of the origin server on behalf of which the request is being made, as described above.) However, a proxy agent GET request takes an entire URL. Therefore, the repeater must construct a proxy request containing the entire URL from the path portion of the URL it receives. Specifically, if the incoming request takes the form:

GET /<origin server>/<path>

then the repeater constructs a proxy request of the form:

GET http://<origin server>/<path>

and if the incoming request takes the form:

GET <origin server>@proxy=<scheme>:<type>@/<path>

then the repeater constructs a proxy request of the form:

GET <scheme>://<origin server>/<path>

Cache Control

[0176] HTTP replies contain directives called cache control directives, which are used to indicate to a cache whether the attached resource may be cached and if so, when it should expire. A Web site administrator configures the Web site to attach appropriate directives. Often, the administrator will not know how long a page will be fresh, and must define a short expiration time to try to prevent caches from serving stale data. In many cases, a Web site operator will indicate a short expiration time only in order to receive the requests (or hits) that would otherwise be masked by the presence of a cache. This is known in the industry as "cache-busting". Although some cache operators may consider cache-busting to be impolite, advertisers who rely on this information may consider it imperative.

[0177] When a resource is stored in a repeater, its cache directives can be ignored by the repeater, because the repeater receives explicit invalidation events to determine when a resource is stale. When a proxy cache is used as the cache at the repeater, the associated cache directives may be temporarily disabled. However, they must be re-enabled when the resource is served from the cache to a client, in order to permit the cache-control policy (including any cache-busting) to take effect.

[0178] The present invention contains mechanisms to prevent the proxy cache within a repeater from honoring cache control directives, while permitting the directives to be served from the repeater.

[0179] When a reflector serves a resource to a repeater in B4, it replaces all cache directives by modified directives that are ignored by the repeater proxy cache. It does this by prefixing a distinctive string such as "wr-" to the beginning of the HTTP tag. Thus, "expires" becomes "wr-expires", and "cache-control" becomes "wr-cache-control". This prevents the proxy cache itself from honoring the directives. When a repeater serves a resource in C4, and the requesting client is not another repeater, it searches for HTTP tags beginning with "wr-" and removes the "wr-". This allows the clients retrieving the resource to honor the directives.

Resource Revalidation

[0180] There are several cases where a resource may be cached so long as the origin server is consulted each time it is served. In one case, the request for the resource is attached to a so-called "cookie". The origin server must be presented with the cookie to record the request and determine whether the cached resource may be served or not In another case, the request for the resource is attached to an authentication header (which identifies the requester with a user id and password). Each new request for the resource must be tested at the origin server to assure that the requester is authorized to access the resource.

[0181] The HTTP 1.1 specification defines a reply header tided "Must-Revalidate" which allows an origin server to instruct a proxy cache to "revalidate" a resource each time a request is received. Normally, this mechanism is used to determine whether a resource is still fresh. In the present invention, Must-Revalidate makes it possible to ask an origin server to validate a request that is otherwise served from a repeater.

[0182] The reflector rule base contains information that determines which resources may be repeated but must be revalidated each time they are served. For each such resource, in B4, the reflector attaches a Must-Revalidate header. Each time a request comes to a repeater for a cached resource marked with a Must-Revalidate header, the request is forwarded to the reflector for validation prior to serving the requested resource.

Cache Quotas

[0183] The cache component of a repeater is shared among those subscribers that reflect clients to that repeater. In order to allow subscribers fair access to storage facilities, the cache may be extended to support

quotas.

[0184] Normally, a proxy cache may be configured with a disk space threshold T. Whenever more than T bytes are stored in the cache, the cache attempts to find resources to eliminate.

[0185] Typically a cache uses the least-recently-used (LRU) algorithm to determine which resources to eliminate; more sophisticated caches use other algorithms. A cache may also support several threshold values-for instance, a lower threshold which, when reached, causes a low priority background process to remove items from the cache, and a higher threshold which, when reached, prevents resources from being cached until sufficient free disk space has been reclaimed.

[0186] If two subscribers A and B share a cache, and more resources of subscriber A are accessed during a period of time than resources of subscriber B, then fewer of B's resources will be in the cache when new requests arrive. It is possible that, due to the behavior of A, B's resources will never be cached when they are requested. In die present invention, this behavior is undesirable. To address this issue, the invention extends the cache at a repeater to support cache quotas.

[0187] The cache records the amount of space used by each subscriber in D_S, and supports a configurable threshold T_S for each subscriber.

[0188] Whenever a resource is added to the cache (at C3), the value D_S is updated for the subscriber providing the resource. If D_S is larger than T_S, the cache attempts to find resources to eliminate, from among those resources associated with subscriber S. The cache is effectively partitioned into separate areas for each subscriber.

[0189] The original threshold T is still supported. If the sum of reserved segments for each subscriber is smaller than the total space reserved in the cache, the remaining area is "common" and subject to competition among subscribers.

[0190] Note, this mechanism might be implemented by modifying the existing proxy cache discussed above, or it might also be implemented without modifying the proxy cache-if the proxy cache at least makes it possible for an external program to obtain a list of resources in the cache, and to remove a given resource from the cache.

Rewriting from Repeaters

[0191] When a repeater receives a request for a resource, its proxy cache may be configured to determine whether a peer cache contains the requested resource. If so, the proxy cache obtains the resource from the peer cache, which can be faster than obtaining it from the origin server (the reflector). However, a consequence of this is that rewritten HTML resources retrieved from the peer cache would identify the wrong repeater. Thus, to allow for cooperating proxy caches, resources are preferably rewritten at the repeater.

[0192] When a resource is rewritten for a repeater, a special tag is placed at the beginning of the resource. When constructing a reply, the repeater inspects the tag to determine whether the resource indicates that additional rewriting is necessary. If so, the repeater modifies the resource by replacing references to the old repeater with references to the new repeater.

[0193] It is only necessary to perform this rewriting when a resource is served to the proxy cache at another repeater.

Repeater-Side Include

[0194] Sometimes, an origin server constructs a custom resource for each request (for instance, when inserting an advertisement based on the history of the requesting client). In such a case, that resource must be served locally rather than repeated. Generally, a custom resource contains, along with the custom information, text and references to other, repeatable, resources.

[0195] The process that assembles a "page" from a text resource and possibly one or more image resources is performed by the Web browser, directed by HTML. However, it is not possible using HTML to cause a browser to assemble a page using text or directives from a separate resource. Therefore, custom resources often necessarily contain large amounts of static text that would otherwise be repeatable.

[0196] To resolve this potential inefficiency, repeaters recognize a special directive called a "repeater side include". This directive makes it possible for the repeater to assemble a custom resource, using a combination of repeatable and local resources. In this way, the static text can be made repeatable, and only the special directive need be served locally by the reflector.

[0197] For example, a resource X might consist of custom directives selecting an advertising banner, followed by a large text article. To make this resource repeatable, the Web site administrator must break out a second resource, Y, to select the banner. Resource X is modified to contain a repeater-side include directive identifying resource Y, along with the article. Resource Y is created and contains only the custom directives selecting an ad banner. Now resource X is repeatable, and only resource Y, which is relatively small, is not repeatable.

[0198] When a repeater constructs a reply, it determines whether the resource being served is an HTML resource, and if so, scans it for repeater-side include directives. Each such directive includes a URL, which the repeater resolves and substitutes in place of the directive. The entire resource must be assembled before it is served, in order to determine its final size, as the size is included in a reply header ahead of the resource.

[0199] Thus, a method and apparatus for dynamically replicating selected resources in computer networks is provided. One skilled in the art will appreciate that the present invention can be practiced by other than the de-

scribed embodiments, which are presented for purposes of illustration and not limitation, and the present invention is limited only by the claims that follow.

Claims

1. A system operative in a computer network (100) in which clients (106) connect to an origin server (102), the system comprising:

   a set of repeater servers (104a ...), distinct from the origin server (102), for hosting at least some of the embedded objects of web pages that are normally hosted by the origin server (102);

   a routine for modifying at least one embedded object URL of a web page to resolve to a server in the set of repeater servers instead of to the origin server;

   a server selector mechanism constructed and adapted to identify, for a particular client request, an appropriate repeater server from the set of repeater servers, wherein in response to requests for the web page, generated by the clients, the web page including the modified embedded object URL is served from the origin server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the server selector mechanism; **characterised in that** the server selector mechanism identifies a repeater server based on the requesting client's location and data relating to network traffic conditions.

2. A system as in claim 1 wherein, if a given one of the repeater servers (104a ...) fails, another repeater server will take over the role of the failed server.

3. A system as in claim 1 wherein the server selecting mechanism also identifies a repeater server based on the load on the content servers.

4. A system as in claim 1 wherein the server selecting mechanism includes a network map for use in directing a request for the embedded object generated by a client.

5. A system as in any one of claims 1-4 wherein the server selector mechanism is co-located with the origin server (102).

6. A method of serving a page supported at an origin server, said method operating a system operative in a computer network, the system comprising: a set of repeater servers (104a, ...), distinct from an origin server (102), for hosting at least some of the embedded objects of web pages that are normally hosted by the origin server; a routine for causing at least one embedded object URL of a web page to resolve to a server in the set of repeater servers instead of to the origin server; and a server selector mechanism constructed and adapted to identify, for a particular client request, an appropriate repeater server from the set of repeater servers, wherein the server selector mechanism identifies a repeater server based on the requesting client's location and data relating to network traffic conditions, the page comprising a base document having associated therewith embedded objects, each embedded object being identified by a uniform resource locator (URL), the method comprising:

   modifying the URL of an embedded object to generate a modified URL, the modified URL including a new hostname referencing a repeater server in the set of repeater servers;

   in response to a request to serve the page, serving the page with the modified URL;

   attempting to serve the embedded object from a repeater server as identified by the new hostname;

   if a copy of the embedded object is unavailable at the repeater server, obtaining the embedded object from the origin server.

7. A method as in claim 6 wherein the modifying of the URL comprises:

   adding the new hostname to the original URL.

8. A method as in claim 7 wherein the original hostname is maintained as part of the modified URL.

9. A method as in claim 6 wherein the original URL includes an origin server name and a path, and wherein the modified URL includes a repeater server name.

10. A method as in claim 9 wherein the original URL has the form:

    "http://<server>/<path>"

    where "<server>" is the origin server name and where "'<path>" is a path to the resource designated by the URL, and wherein the modified URL has the form:

    "http://<repeater>/<server>/<path>"

    where "<repeater>" is the hostname of the repeater server.

11. A method as in claim 9 wherein the original URL

has the form:

"http://<server>/<path>"

where "<server>" is the origin server name and where "<path>" is a path to the resource designated by the URL, and wherein the modified URL has the form:

"http://<repeater>/<server>/<path>"

where "<repeater>" is a hostname of a set of repeater servers.

12. A method of serving a page as in claim 7, wherein

the modifying of the URL comprises: modifying the URL for the embedded page object to include a hostname prepended to an origin server-supplied domain name and path;

the serving of the page comprises: serving the page with the modified URL from the origin server to a client's browser;

the method further comprising:

returning to the client's browser an IP address of the identified repeater server to enable the browser to attempt to retrieve the object from that repeater server.

13. A method in claim 12 wherein copies of the embedded page object are stored on a subset of the set of repeater servers.

14. A method as in claim 6 wherein the modifying of the embedded object URL causes the address of the embedded object in the page to resolve to an address other than an origin server address by adding given data to a origin server-supplied address to generate an alternate address.

15. A method as in claim 14 wherein the computer network is the Internet and the address of the object is a uniform resource locator (URL).

16. A method as in claim 15 wherein the alternate address is formed by prepending the given data to a part of the origin server-supplied address.

17. A method as in any one of claims 14-16 wherein the resolving of the alternate address comprises:

using a requesting client's location and data relating to Internet traffic conditions to identify the repeater server.

18. A method as in claim 6, the method, further comprising:

replicating a set of page objects across a wide area network of repeater servers.

19. The method as described in claim 6, further comprising:

for each modified embedded object URL, identifying at least one repeater server from which the embedded object may be retrieved.

20. The method as in claim 19 wherein the identifying comprises:

resolving a request for an embedded object as a function of a requesting client's location.

21. The method as in claim 19 wherein the identifying comprises:

resolving a request for an embedded object as a function of a requesting client's location and then-current Internet traffic conditions.

22. A method as in claim 6, further comprising:

responsive to a request for the embedded object, resolving the hostname to an IP address of a particular repeater server.

23. A method as in claim 22 wherein the resolving of the hostname includes:

identifying a subset of repeater servers (104a, ...) that may be able to serve the embedded object; and identifying the particular repeater server from the identified subset of repeater servers.

24. A method as in claim 23 wherein the identifying of the particular repeater server is based on a location of the client machine and on Internet traffic conditions

25. A method as in claim 6 further comprising:

distributing a set of page objects across a network of repeater servers (104a, ...), wherein the network of repeater servers are organized into a set of regions; in response to a client request for an embedded object of a page:

resolving the client request as a function of a location of the client and Internet traffic conditions to identify a given region;

returning to the client an IP address of a particular one of the repeater servers within the given region.

26. A system as in claim 1 wherein the routine for modifying modifies a URL to include a hostname prepended to a domain name and path.

27. A system as in claim 1 wherein the routine for modifying adds a new hostname to the original URL.

28. A system as in claim 1 further comprising:

at least one DNS server for providing name-to-address resolution of hostnames in embedded object URLs.

Patentansprüche

1. Ein in einem Computernetzwerk (100) betriebenes System, in welchem Clients (106) mit einem Ausgangsserver (102) Verbindung aufnehmen, wobei das System beinhaltet:

eine Serie von Repeaterservern (104a ...), die vom Ausgangsserver (102) verschieden sind und als Host für mindestens einige der eingebetteten Objekte von Webseiten dienen, die normalerweise vom Ausgangsserver (102) als Host verwaltet werden,

eine Routine, die mindestens eine URL eines eingebetteten Objekts einer Webseite so modifiziert, daß sie zu einem Server in der Serie von Repeaterservern anstatt zum Ausgangsserver aufgelöst wird,

einen Server-Auswahlmechanismus, der so aufgebaut und angeordnet ist, daß er auf eine bestimmte Client-Anforderung hin einen geeigneten Repeaterserver aus der Serie von Repeaterservern identifiziert, wobei als Reaktion auf von den Clients erzeugte Anforderungen nach der Webseite vom Ausgangsserver die mit der modifizierten URL des eingebetteten Objekts versehene Webseite bereitgestellt wird und das eingebettete Objekt, das durch die modifizierte URL des eingebetteten Objekts identifiziert wird, von einem gegebenen der vom Server-Auswahlmechanismus identifizierten Repeaterserver bereitgestellt wird, **dadurch gekennzeichnet, daß** der Server-Auswahlmechanismus einen Repeaterserver anhand des Standortes des anfordernden Clients und anhand von Daten im Zusammenhang mit dem Datenverkehrsaufkommen im Netzwerk identifiziert.

2. System nach Anspruch 1, in welchem bei einem Ausfall eines der Repeater-server (104a...) ein anderer Repeaterserver die Rolle des ausgefallenen Servers übernimmt.

3. System nach Anspruch 1, in welchem der Server-Auswahlmechanismus außerdem einen Repeaterserver in Abhängigkeit von der Auslastung der Contentserver identifiziert.

4. System nach Anspruch 1, in welchem der Server-Auswahlmechanismus eine Netzwerkübersicht beinhaltet, anhand derer eine von einem Client erzeugte Anforderung nach dem eingebetteten Objekt zielgerichtet versandt werden kann.

5. System nach einem oder mehreren der Ansprüche 1 bis 4, in welchem der Server-Auswahlmechanismus am selben Standort angeordnet ist wie der Ausgangsserver (102).

6. Ein Verfahren zum Bereitstellen einer auf einem Ausgangsserver unterstützten Seite, wobei das genannte Verfahren zur Nutzung eines in einem Computernetzwerk betriebenen Systems dient und das System folgendes beinhaltet: Eine Serie von Repeaterservern (104a...), die vom Ausgangsserver (102) verschieden sind und als Hostsystem für mindestens einige der eingebetteten Objekte von Webseiten dienen, die normalerweise vom Ausgangsserver als Host verwaltet werden, ferner eine Routine, die bewirkt, daß mindestens eine URL eines eingebetteten Objekts einer Webseite zu einem Server in der Serie von Repeaterservern anstatt zum Ausgangsserver aufgelöst wird, des weiteren einen Server-Auswahlmechanismus, der so aufgebaut und angeordnet ist, daß er auf eine bestimmte Client-Anforderung hin einen geeigneten Repeaterserver aus der Serie von Repeaterservern identifiziert, wobei der Server-Auswahlmechanismus einen Repeaterserver anhand des Standortes des anfordernden Clients und anhand von Daten im Zusammenhang mit dem Datenverkehrsaufkommen im Netzwerk identifiziert, wobei die Seite ein Basisdokument beinhaltet, welchem eingebettete Objekte zugeordnet sind, von denen jedes durch eine URL (Uniform Resource Locator) identifiziert wird, wobei das Verfahren folgendes beinhaltet:

Modifizieren der URL eines eingebetteten Objekts zur Erzeugung einer modifizierten URL mit einem zusätzlichen neuen Hostnamen, der auf einen Repeaterserver in der Serie der Repeaterserver verweist;

Bereitstellen der Seite mit der modifizierten URL als Reaktion auf eine Anforderung zum Bereitstellen der Seite;

Versuch zum Bereitstellen des eingebetteten Objekts von einem durch den neuen Host-Na-

men identifizierten Repeaterserver aus;

Abrufen des eingebetteten Objekts vom Ausgangsserver, wenn eine Kopie des eingebetteten Objekts auf dem Repeaterserver nicht verfügbar ist.

7. Verfahren nach Anspruch 6, wobei das Modifizieren der URL
das Hinzufügen des neuen Hostnamens zur ursprünglichen URL beinhaltet.

8. Verfahren nach Anspruch 7, wobei der ursprüngliche Hostname als Bestandteil der modifizierten URL beibehalten wird.

9. Verfahren nach Anspruch 6, wobei die ursprüngliche URL einen Ausgangs-servernamen und einen Pfad beinhaltet, und wobei die modifizierte URL einen Repeaterservernamen beinhaltet.

10. Verfahren nach Anspruch 9, wobei die ursprüngliche URL die Form
"http://<server>/<pfad>"
hat, wobei "<server>" der Ausgangsservername ist, wobei "<pfad>" ein Pfad ist, der zu der durch die URL angegebenen Ressource führt, und wobei die modifizierte URL die Form
"http://<repeater>/<server>/<pfad>"
hat, wobei "<repeater>" der Hostname des Repeaterservers ist.

11. Verfahren nach Anspruch 9, wobei die ursprüngliche URL die Form
"http://<server>/<pfad>"
hat, wobei "<server>" der Ausgangsservername ist, wobei "<pfad>" ein Pfad ist, der zu der durch die URL angegebenen Ressource führt, und wobei die modifizierte URL die Form
"http://<repeater>/<server>/<pfad>"
hat, wobei "<repeater>" ein Hostname einer Serie von Repeaterservern ist.

12. Verfahren nach Anspruch 7, wobei
das Modifizieren der URL die Schritte beinhaltet, die URL für das eingebettete Seitenobjekt so zu modifizieren, daß sie einen Hostnamen enthält, der einem vom Ausgangsserver bereitgestellten Domainnamen und Pfad vorangestellt ist,
das Bereitstellen der Seite das Bereitstellen der Seite mit der modifizierten URL vom Ausgangsserver gegenüber dem Browser eines Clients beinhaltet,
das Verfahren weiterhin folgendes beinhaltet:

Zurücksenden einer IP-Adresse des identifizierten Repeaterservers an den Browser des Clients, so daß der Browser in die Lage versetzt wird, einen Versuch zum Abrufen des Objekts von diesem Repeaterserver zu unternehmen.

13. Verfahren nach Anspruch 12, wobei Kopien des eingebetteten Seitenobjekts auf einer Untermenge der Serie von Repeaterservern gespeichert sind.

14. Verfahren nach Anspruch 6, wobei das Modifizieren der URL des einge-betteten Objekts bewirkt, daß die Adresse des eingebetteten Objekts auf der Seite zu einer anderen Adresse als einer Ausgangsserver-Adresse aufgelöst wird, indem einer vom Ausgangsserver bereitgestellten Adresse vorgegebene Daten hinzugefügt werden, um so eine Alternativadresse zu erzeugen.

15. Verfahren nach Anspruch 14, wobei das Computernetzwerk das Internet ist und die Adresse des Objekts eine URL (Uniform Resource Locator) ist.

16. Verfahren nach Anspruch 15, wobei die Alternativadresse gebildet wird, indem die gegebenen Daten einem Teil der vom Ausgangsserver bereitgestellten Adresse vorangestellt werden.

17. Verfahren nach einem oder mehreren der Ansprüche 14-16, wobei das Auflösen der Alternativadresse den Schritt beinhaltet,
unter Verwendung des Standortes eines anfordernden Clients sowie von Daten im Zusammenhang mit dem Datenverkehrsaufkommen im Internet den Repeater-server zu identifizieren.

18. Verfahren nach Anspruch 6, wobei das Verfahren weiterhin

das Nachbilden einer Serie von Seitenobjekten über ein Weitverkehrsnetz von Repeaterservern beinhaltet.

19. Verfahren nach Anspruch 6 beschrieben, wobei das Verfahren weiterhin
für jede modifizierte URL eines eingebetteten Objekts das Identifizieren mindestens eines Repeaterservers beinhaltet, von welchem aus das eingebettete Objekt abgerufen werden kann.

20. Verfahren nach Anspruch 19, wobei das Identifizieren
das Auflösen einer Anforderung nach einem eingebetteten Objekt als Funktion des Standortes eines anfordernden Clients beinhaltet.

21. Verfahren nach Anspruch 19, wobei das Identifizieren
das Auflösen einer Anforderung nach einem eingebetteten Objekt als Funktion des Standortes eines anfordernden Clients sowie des dann aktuellen In-

ternet-Datenverkehrsaufkommens beinhaltet.

22. Verfahren nach Anspruch 6, wobei das Verfahren weiterhin
als Reaktion auf eine Anforderung nach dem eingebetteten Objekt das Auflösen des Hostnamens zu einer IP-Adresse eines bestimmten Repeaterservers beinhaltet.

23. Verfahren nach Anspruch 22, wobei das Auflösen des Hostnamens
das Identifizieren einer Untermenge von Repeaterservern (104a, ...), die in der Lage sein können, das eingebettete Objekt bereitzustellen, sowie das Identifizieren des bestimmten Repeaterservers aus der identifizierten Untermenge von Repeaterservern beinhaltet.

24. Verfahren nach Anspruch 23, wobei das Identifizieren des bestimmten Repeaterservers anhand des Standorts der Client-Maschine und des Internet-Datenverkehrsaufkommens erfolgt.

25. Verfahren nach Anspruch 6, wobei das Verfahren weiterhin
das Verteilen einer Serie von Seitenobjekten über ein Netzwerk von Repeaterservern (104a, ...), wobei das Netzwerk von Repeaterservern als Serie von Regionen organisiert ist, beinhaltet, als Reaktion auf eine Client-Anforderung nach einem eingebetteten Objekt einer Seite durch
Auflösen der Client-Anforderung als Funktion eines Standortes des Clients sowie des Internet-Datenverkehrsaufkommens zum Identifizieren einer gegebenen Region und
Zurücksenden einer IP-Adresse eines bestimmten Repeaterservers innerhalb der gegebenen Region an den Client.

26. System nach Anspruch 1, in welchem die Routine zum Modifizieren eine URL so modifiziert, daß sie einen Hostnamen enthält, der einem vom Ausgangs-server bereitgestellten Domainnamen und Pfad vorangestellt ist.

27. System nach Anspruch 1, in welchem die Routine zum Modifizieren der ursprünglichen URL einen neuen Hostnamen hinzufügt.

28. System nach Anspruch 1, ferner beinhaltend mindestens einen DNS-Server zur Bereitstellung einer Namens-Adressen-Auflösung von Hostnamen in den URLs in eingebetteten Objekten.

Revendications

1. Système opérationnel dans un réseau informatique (100) dans lequel des clients (106) se connectent à un serveur d'origine (102), le système comprenant:

un ensemble de serveurs répéteurs (104a ...), distincts du serveur d'origine (102), pour héberger au moins certains des objets intégrés de pages de la toile mondiale qui sont normalement hébergés par le serveur d'origine (102),

une routine pour modifier au moins une adresse URL d'objet intégré d'une page de la toile mondiale à définir dans un serveur de l'ensemble de serveurs répéteurs à la place du serveur d'origine,

un mécanisme sélecteur de serveur conçu et adapté pour identifier, pour une demande de client particulier, un serveur répéteur approprié parmi l'ensemble de serveurs répéteurs, dans lequel en réponse à des demander de la page de la toile mondiale, générées par les clients, la page de la toile mondiale incluant l'adresse URL d'objet intégré modifiée est mise en service depuis le serveur d'origine et l'objet intégré identifié par l'adresse URL d'objet intégré modifiée est mis en service depuis un serveur donné parmi les serveurs répéteurs lorsqu'identifié par le mécanisme sélecteur de serveur, **caractérisé en ce que** le mécanisme sélecteur de serveur identifie un serveur répéteur sur la base de l'emplacement du client demandeur et des données relatives aux conditions de trafic du réseau.

2. Système selon la revendication 1, dans lequel, si un serveur donné parmi les serveurs répéteurs (104a ...) est défectueux, un autre serveur répéteur va reprendre le rôle du serveur défectueux.

3. Système selon la revendication 1, dans lequel le mécanisme de sélection de serveur identifie également un serveur répéteur sur la base de la charge des serveurs de contenus.

4. Système selon la revendication 1, dans lequel le mécanisme de sélection de serveur inclut une carte du réseau à utiliser pour adresser une demande concernant l'objet intégré générée par un client.

5. Système selon l'une quelconque des revendications 1 à 4, dans lequel le mécanisme sélecteur de serveur est situé conjointement avec le serveur d'origine (102).

6. Procédé pour mettre en service une page supportée au niveau d'un serveur d'origine, ledit procédé utilisant un système opérationnel dans un réseau informatique, le système comportant : un ensemble

de serveurs répéteurs (104, ...), distincts d'un serveur d'origine (102), pour héberger au moins certains des objets intégrés de pages de la toile mondiale qui sont normalement héhergés par le serveur d'origine, une routine pour amener au moins une adresse URL d'objet intégré d'une page de la toile mondiale à être définie dans un serveur de l'ensemble de serveurs répéteurs à la place du serveur d'origine, et un mécanisme sélecteur de serveur conçu et adapté pour identifier, pour une demande de client particulier, un serveur répéteur approprié parmi l'ensemble de serveurs répéteurs, dans lequel le mécanisme sélecteur de serveur identifie un serveur répéteur sur la base de l'emplacement du client demandeur et de données relatives aux conditions de trafic du réseau, la page comportant un document de base ayant des objets intégrés qui lui sont associés, chaque objet intégré étant identifié par une adresse uniformisée de ressources (URL), le procédé comprenant:

modifier l'adresse URL d'un objet intégré pour générer une adresse URL modifiée, l'adresse URL modifiée incluant un nouveau nom d'hôte faisant référence à un serveur répéteur de l'ensemble de serveurs répéteurs,

en réponse à une demande de mise en service de la page, mettre en service la page avec l'adresse URL modifiée,

tenter de mettre en service l'objet intégré à partir d'un serveur répéteur lorsqu'identifié par le nouveau nom d'hôte.

si une copie de l'objet intégré est indisponible dans le serveur répéteur, obtenir l'objet intégré depuis le serveur d'origine.

7. Procédé selon la revendication 6, dans lequel la modification de l'adresse URL comporte:

l'ajout du nouveau nom d'hôte à l'adresse URL d'origine.

8. Procédé selon la revendication 7,'dans lequel le nom d'hôte d'origine est maintenu comme un partie de l'adresse URL modifiée.

9. Procédé selon is revendication 6, dans lequel l'adresse URL d'origine inclut un nom de serveur d'origine et un chemin, et dans lequel l'adresse URL modifiée inclut un nom de serveur répéteur.

10. Procédé selon la revendication 9, dans lequel l'adresse URL d'origine a la forme:
"http://<serveur>/<chemin>"
où "<serveur>" est le nom de serveur d'origine et où "<chemin>" est un chemin jusqu'à la ressource désignée par l'adresse URL, et dans lequel l'adresse URL modifiée a la forme:
"http://<répéteur>/<serveur>/<chemin>"
où "<répéteur>" est le nom d'hôte du serveur répéteur.

11. Procédé selon la revendication 9, dans lequel l'adresse URL d'origine a la forme :
"http://<serveur>/<chemin>"
où "<serveur>" est le nom de serveur d'origine et où "<chemin>" est un chemin jusqu'à la ressource désignée par l'adresse URL, et dans lequel l'adresse URL modifiée a la forme:
"http://<répéteur>/<serveur>/<chemin>"
où "<répéteur>" est un nom d'hôte d'un ensemble de serveurs répéteurs.

12. Procédé pour mettre en service une page selon la revendication 7, dans lequel
la modification de l'adresse URL comporte: la modification de l'adresse URL pour l'objet de page intégré afin d'inclure un nom d'hôte ajouté en préfixe à un nom de domaine et chemin fournis par le serveur d'origine,
la mise en service de la page comporte: la mise en service de la page avec l'adresse URL modifiée depuis le serveur d'origine vers un navigateur de client,
le pracédé comprenant en outre l'étape consistant à :

retourner au navigateur de client une adresse de protocole Internet (IP) du serveur répéteur identifié pour permettre au navigateur de tenter de récupérer l'objet à partir de ce serveur répéteur.

13. Procédé selon la revendication 12, dans lequel des copies de l'objet de page intégré sont mémorisées sur un sous-ensemble de l'ensemble de serveurs répéteurs.

14. Procédé selon la revendication 6, dans lequel la modification de l'adresse URL d'objet intégré amène l'adresse de l'objet intégré dans la page à être définie à une adresse autre qu'une adresse de serveur d'origine en ajoutant des données données à une adresse fournie par le serveur d'origine afin de générer une autre adresse.

15. Pracédé selon la revendication 14, dans lequel le réseau informatique est l'Internet et l'adresse de l'objet est une adresse uniformisée de ressources (URL).

16. Pracédé selon la revendication 15, dans lequel l'autre adresse est formée en ajoutant en préfixe les

données données à une partie de l'adresse fournie par le serveur d'origine.

17. procédé selon l'une quelconque des revendications 14 à 16, dans lequel la définition de l'autre adresse comporte :

    l'utilisation de l'emplacement du client demandeur et des données relatives aux conditions de trafic de l'Internet pour identifier le serveur répéteur.

18. Procédé selon la revendication 6, le procédé comprenant en outre :

    la réplication d'un ensemble d'objets de page à travers un réseau longue distance de serveurs répéteurs.

19. Procédé selon la revendication 6, comprenant en outre :

    pour chaque adresse URL d'abjet intégré modifiée, l'identification d'au moins un serveur répéteur à partir duquel l'objet intégré peut être récupéré.

20. Procédé selon la revendication 19, dans lequel l'identification comprend :

    la définition d'une demande pour un objet intégré en fonction de l'emplacement du client demandeur.

21. Pracédé selon la revendication 19, dans lequel l'identification comprend:

    la définition d'une demande d'un objet intégré en fonction de l'emplacement du client demandeur et des conditions de trafic de l'internet courantes à ce moment.

22. Procédé selon la revendication 6, comprenant en outre :

    en réponse à une demande de l'objet intégré, la définition du nom d'hôte par une adresse IP d'un serveur répéteur particulier.

23. Procédé selon la revendication 22, dans lequel la définition du nom d'hôte inclut :

    l'identification d'un sons-ensemble de serveurs répéteurs (104a, ...) qui peuvent être capables de mettre en service l'objet intégré, et l'identification du serveur répéteur particulier à partir du sous-ensemble de serveurs répéteurs identifiés.

24. Procédé selon la revendication 23, dans lequel l'identification du serveur répéteur particulier est basée sur un emplacement de la machine cliente et sur les conditions de trafic de l'Internet.

25. Procédé selon la revendication 6, comprenant en outre :

    répartir un ensemble d'objets de page à travers un réseau de serveurs répéteurs (104a, ...), dans lequel le réseau de serveurs répéteurs est organisé en un ensemble de régions, en réponse à une demande de client concernant un objet intégré d'une page :

        définir la demande de client en fonction d'un emplacement du client et des conditions de trafic de l'Internet pour identifier une région donnée,

        renvoyer au client une adresse IP d'un serveur particulier parmi les serveurs répéteurs situés dans la région donnée.

26. Système selon la revendication 1, dans lequel la routine de modification modifie une adresse URL pour inclure un nom d'hôte ajouté en préfixe à un nom de domaine et chemin.

27. Système selon la revendication 1, dans lequel la routine de modification ajoute un nouveau nom d'hôte à l'adresse URL d'origine.

28. Système selon la revendication 1, comprenant en outre :

    au moins un serveur de noms de domaine (DNS) pour fournir une définition nom-adresse de noms d'hôte dans des adresses URL d'objet intégré.

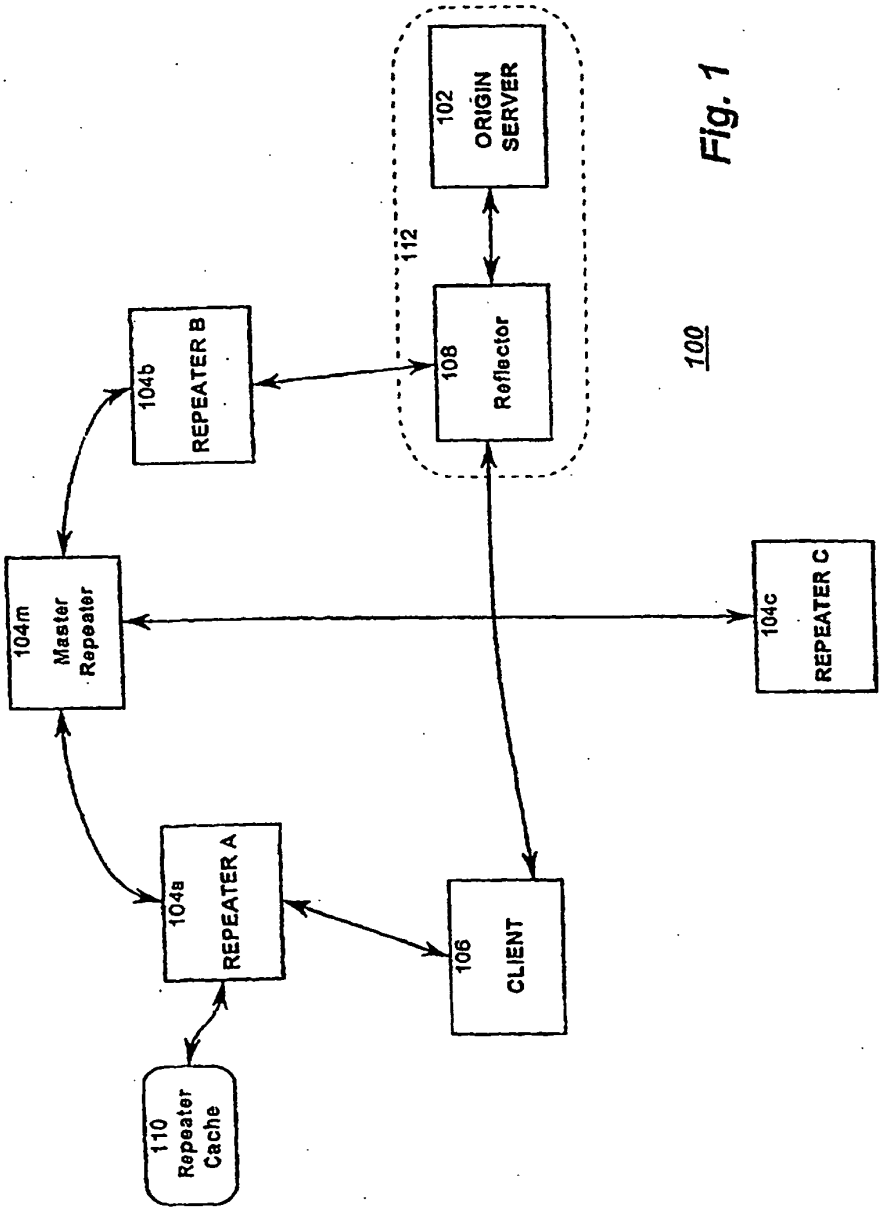

102
ORIGIN SERVER
112
108
Reflector
104b
REPEATER B
104m
Master Repeater
104c
REPEATER C
104a
REPEATER A
106
CLIENT
110
Repeater Cache

Fig. 1

100

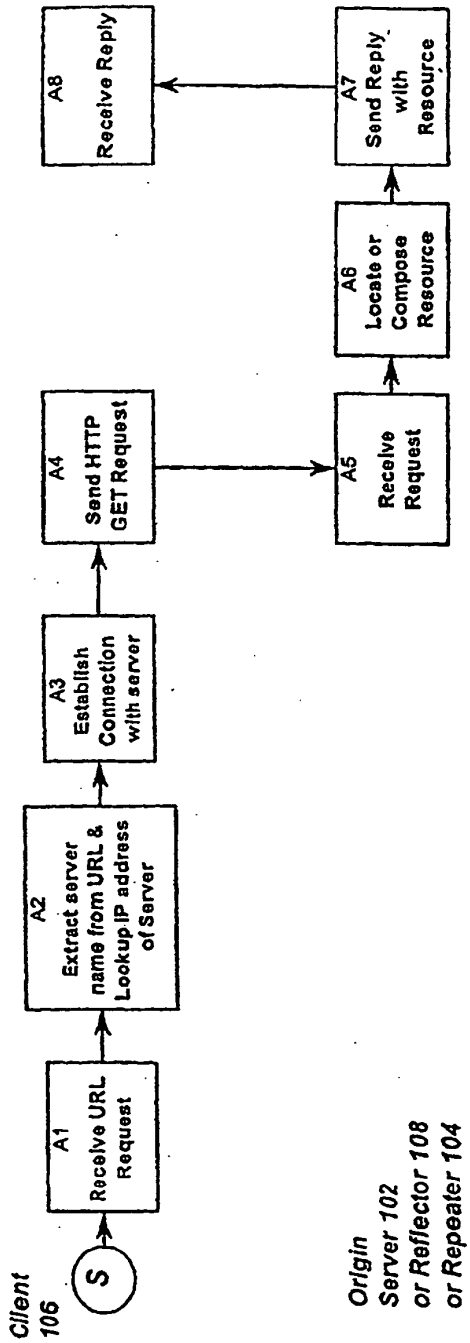

A1
Receive URL Request
A2
Extract server name from URL & Lookup IP address of Server
A3
Establish Connection with server
A4
Send HTTP GET Request
A5
Receive Request
A6
Locate or Compose Resource
A7
Send Reply with Resource
A8
Receive Reply
S
Client 106
Origin Server 102 or Reflector 108 or Repeater 104

Fig. 2

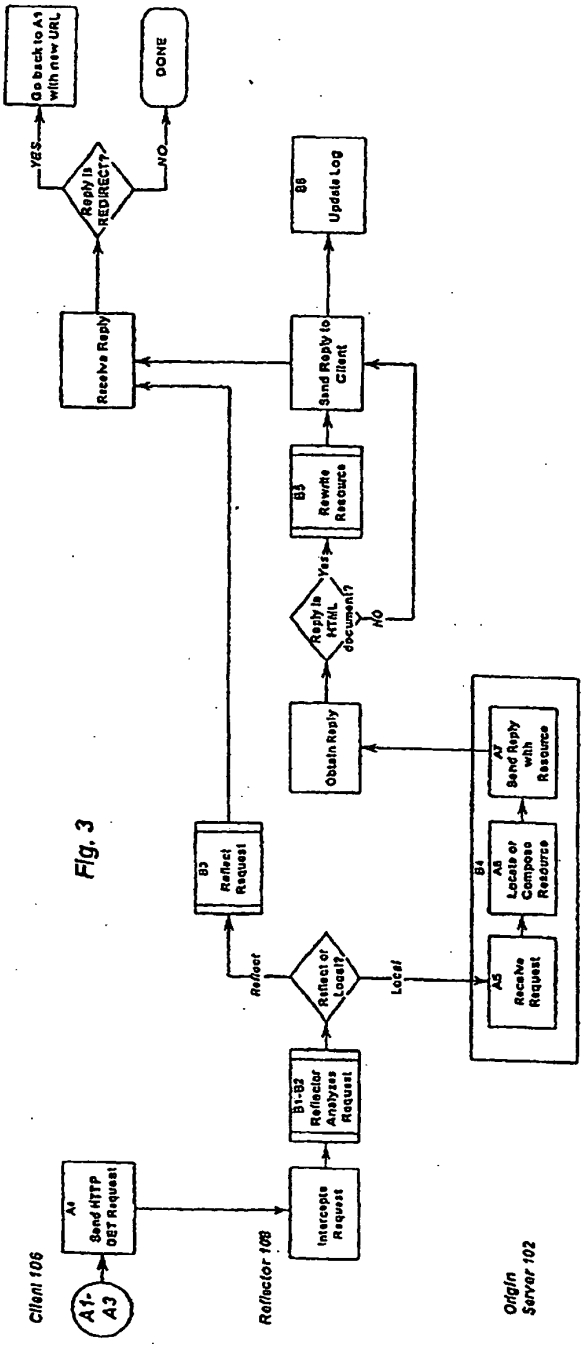

Client 106
Reflector 108
Origin Server 102
A1- A3
A1 Send HTTP GET Request
Intercept Request
B1-B2 Reflector Analyzes Request
Reflect or Local?
Reflect
Local
B3 Reflect Request
A5 Receive Request
B4
A6 Locate or Compose Resource
A7 Send Reply with Resource
Obtain Reply
Reply is HTML document?
Yes
No
B5 Rewrite Resource
Send Reply to Client
B6 Update Log
Receive Reply
Reply is REDIRECT?
YES
NO
Go back to A1 with new URL
DONE

Fig. 3

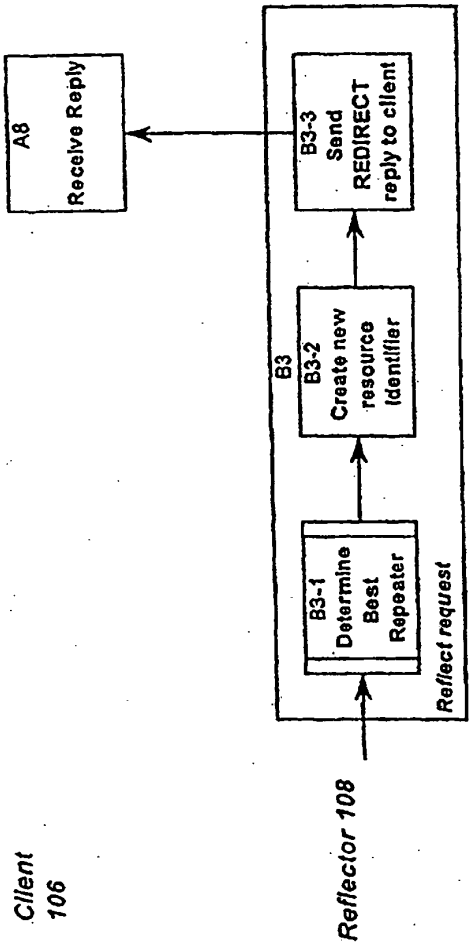

A8
Receive Reply
B3
B3-3
Send
REDIRECT
reply to client
B3-2
Create new
resource
identifier
B3-1
Determine
Best
Repeater
Reflect request
Reflector 108
Client
106

Fig. 4

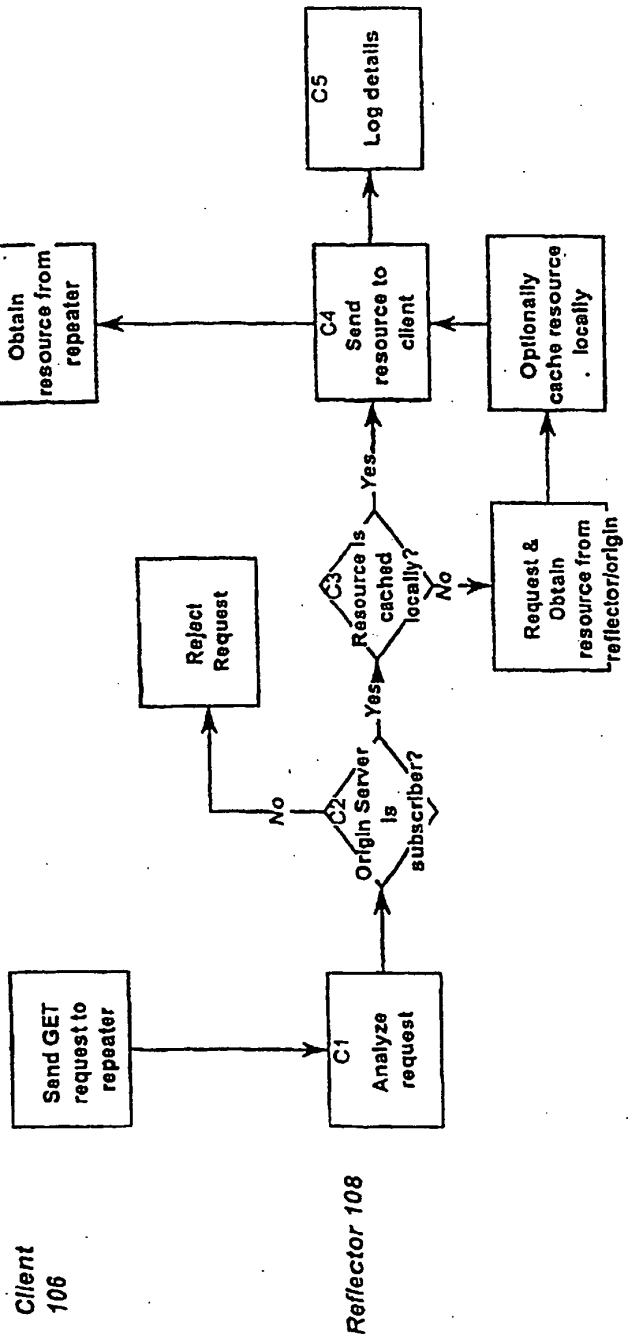

Client 106
Reflector 108
Send GET request to repeater
C1 Analyze request
C2 Origin Server is subscriber?
No
Yes
Reject Request
C3 Resource is cached locally?
Yes
No
Request & Obtain resource from reflector/origin
Optionally cache resource locally
C4 Send resource to client
Obtain resource from repeater
C5 Log details

Fig. 5

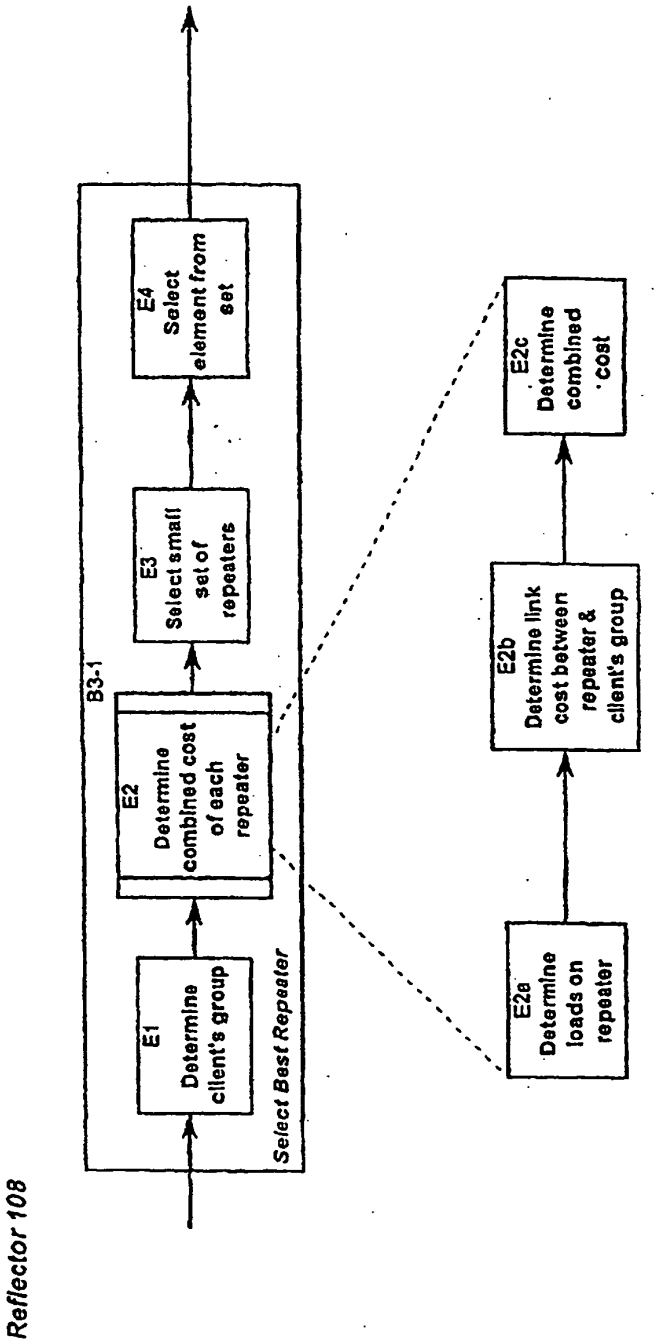

Reflector 108
Select Best Repeater
B3-1
E1
Determine client's group
E2
Determine combined cost of each repeater
E3
Select small set of repeaters
E4
Select element from set
E2a
Determine loads on repeater
E2b
Determine link cost between repeater & client's group
E2c
Determine combined cost

Fig. 6

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2442
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	1-1	US-6181690	2001/01	Civanlar, Mehmet Reha
	1-2	US-6222856	2001/04	Krishnan et al.
	1-3	US-6452925	2002/09	Sistanizadeh et al.
	1-4	US-6779031	2004/08	Picher-Dempsey, Heidi
	1-5			
	1-6			
	1-7			
	1-8			
	1-9			
	1-10			
	1-11			
	1-12			
	1-13			
	1-14			
	1-15			
	1-16			
	1-17			
	1-18			
	1-19			
	1-20			
	1-21			
	1-22			
	1-23			
	1-24			
	1-25			
	1-26			

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2442
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES			
Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	2-1	"Opinion and Order," (Markman – claim construction Order), U.S. Dist. Judge Mark S. DAVIS, Dec. 10, 2008 in Civil Action No. 2:07cv589 (WDK-FBS), in US District Court, Norfolk, VA. [51 pgs.].	
	2-2	Jury VERDICT FORM in Civil Action No. 2:07cv589 (MSD-FBS), January 23, 2009 [3 pgs.]	
	2-3	Non-Final Office Action in U.S. Appln. No. 11/980,689, May 20, 2009.	
	2-4		
	2-5		
	2-6		
	2-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Electronic Patent Application Fee Transmittal

Application Number:	11065412
Filing Date:	23-Feb-2005
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Filer:	Brian Siritzky
Attorney Docket Number:	2711-0073 (0090-US-C3)

Filed as Large Entity

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Extension - 3 months with $0 paid	1253	1	1110	1110

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
Submission- Information Disclosure Stmt	1806	1	180	180
		Total in USD ($)		**1290**

Electronic Acknowledgement Receipt	
EFS ID:	5669585
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	75948
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073 (0090-US-C3)
Receipt Date:	09-JUL-2009
Filing Date:	23-FEB-2005
Time Stamp:	14:22:08
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$1290
RAM confirmation Number	428
Deposit Account	501860
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows: Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees) Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)	

File Listing:					
Document Number	**Document Description**	**File Name**	**File Size(Bytes)/ Message Digest**	**Multi Part /.zip**	**Pages (if appl.)**
1	Extension of Time	PA_EOT_filed_07_09_09.pdf	122728 cb1b7f50089dfebea9cd249fa8a1355f96968c29	no	1
Warnings:					
Information:					
2	Transmittal Letter	PA_IDS_filed_07_09_09.pdf	581827 475da35c59dfd7c3983410f3af6caa61a4dec2ef	no	2
Warnings:					
Information:					
3	Information Disclosure Statement (IDS) Filed (SB/08)	PA_1449s_filed_07_09_09.pdf	276209 9d17e4f74088394033cf63a18e5982c33bb1a551	no	2
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
4	NPL Documents	L3_Jury_Verdict.pdf	309987 543b1fef4304173a9b21497af5f150176d37109f	no	3
Warnings:					
Information:					
5	NPL Documents	O_MarkmanRuling_11_12_08.pdf	8453195 ab691075071cdbe4b6b9bcc2dd0e15ad7e861cb1	no	51
Warnings:					
Information:					
6	Amendment/Req. Reconsideration-After Non-Final Reject	PA_Response_07_09_09.pdf	5143576 ebb88e5076cfe025eea804d89905363 4f425d1f4	no	25
Warnings:					
Information:					
7	Fee Worksheet (PTO-875)	fee-info.pdf	32509 c63d9450946b678576a8b8d959a38131b7e20640	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	14920031		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2442
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: July 9, 2009

PETITION FOR EXTENSION OF TIME

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Petition is hereby made to extend the original due date to cover the date this response/paper is filed for which the requisite fee is:

1. Extension fee required:

Extended	**Large/Small Entity**	
3 months	$1110	$1110
Less extension fees paid since last Action		(-)
TOTAL FEE OF:		**$1110**

Please charge the present and any missing or inadequate fee for this petition to our Deposit Account/Order Nos./ shown below.

Deposit Account No. 501860, Our Order No. 2711-0073

Customer No. 75948
Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

By: /Brian Siritzky/Reg. No. 37,497

Atty: Brian Siritzky, Ph.D., Reg. No.: 37,497

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:		Attorney Docket:	2711-0073
FARBER, David et al.		Group Art Unit:	2442
Application Serial No.:	11/065,412	Examiner:	BLAIR, Douglas B.
Application Filing Date:	February 23, 2005	Confirmation No.:	3296
Title:	**SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date:	July 9, 2009

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference listed on the attached PTO-1449. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This IDS is intended to be in full compliance with the rules, but should the Examiner find any part of its requirement content to have been omitted, prompt notice to that effect is earnestly solicited, along with additional time under Rule 97(f), to enable Applicant to comply fully.

This Information Disclosure Statement is being filed more than three (3) months after the U.S. filing date AND after the mailing date of the first Office Action on the merits, but before the mailing date of a Final Rejection or Notice of Allowance. Please credit or debit Deposit Account No. 501860 the amount required under 37 C.F.R. § 1.17(p) as needed to ensure consideration of the disclosed information.

This IDS cites documents from Civil Action No. 2:07cv589 (MSD-FBS) in related patents, as well as an office action from a related U.S. Patent Application (Non-Final Office Action in U.S. Appln. No. 11/980,689, May 20, 2009.) As this patent application is stored electronically at the PTO, no copies of the action are being provided herewith. If the Examiner requires copies of any of these applications or any additional information regarding any of the documents cited herein, the Examiner is respectfully requested to contact the undersigned at the number provided.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By: **/Brian Siritzky/Reg. No. 37,497** Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 • FAX: (703) 894-6430	

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: SHARED CONTENT DELIVERY INFRASTRUCTURE	Date: July 9, 2009

RESPONSE

Box Amendment
Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicant hereby responds to the non-final Office Action mailed January 9, 2009.

Claim Amendments begin on page 2 hereof.

Remarks begin on page 11 hereof.

Submitted herewith are:

- A **Petition for Extension of Time**, along with the required fee.
- An **Information Disclosure Statement** along with the required fee.

IN THE CLAIMS

Please amend the claims as follows:

1-26. (Canceled)

27. (Currently Amended) A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;

a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level domain name service (DNS) resolution;

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and **[[**the**]]** an embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

28. (Previously Presented) The distributed hosting framework as recited in claim 27, further including a redundant second level name server.

29. (Currently amended) A ~~distributed hosting framework~~ content delivery system operative in a computer network for delivering resources associated with a

plurality of content providers to multiple client machines, the ~~framework~~ system comprising:

at least one shared ~~content~~ repeater server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers; ~~and~~

~~at least one name server that provides domain name service (DNS) resolution,~~

wherein ~~resources~~ at least a first resource associated with the first content provider ~~are~~ is associated with a first alias name, said first alias name being associated with said at least one shared repeater server; and

wherein, ~~that causes the at least one name server to identify the at least one shared content server to serve the resources in response to~~ requests for the first resource ~~resources associated with the first content provider~~ made from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name; and

wherein ~~the resources~~ at least a second resource associated with the second content provider ~~are~~ is associated with a second alias name, said second alias name being associated with said at least one shared repeater server; and wherein

~~that causes the at least one name server to identify the at least one shared content server to serve the resources in response to~~ requests for the ~~resources~~ second resource ~~associated with the second content provider~~ made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name.

30. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 29, wherein the least one shared ~~content~~ repeater server is addressable using the first alias name and the second alias name.

31. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 29, further comprising:

a reflector constructed and adapted to associate the resources with the first alias name based, at least in part, on a relative cost of transmitting the resources to the client machines.

32. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 29 further comprising:

a set of shared ~~content~~ repeater servers that comprises the at least one ~~content~~ repeater server; and

a reflector constructed and adapted to associate ~~a particular~~ the first resource with the first alias name after selecting the at least one shared ~~content~~ repeater server from the set of shared ~~content~~ repeater servers in response to a request for the ~~particular~~ first resource issued by a particular client machine.

33. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 29, wherein at least one of the first resource and the second resource ~~resources~~ is an embedded object.

34. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 29, wherein the requests comprise host tag names that identify specific origin servers from which the resources originate.

35. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 34, wherein the host tag names are included in HTTP headers.

36. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 32, further comprising:

a load determination mechanism constructed and adapted to characterize load on the set of ~~content~~ repeater servers for use by the reflector in selecting the at least one shared ~~content~~ repeater server.

37. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 36, wherein the reflector comprises the load mechanism.

38. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 37, wherein the reflector is co-located with an origin server at which the first resource originates.

39. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 36, further comprising:

a cost determination mechanism constructed and adapted to measure a cost of delivering resources from each of the set of shared ~~content~~ repeater servers to the particular client machine, the cost determination mechanism for use by the reflector in selecting the at least one shared ~~content~~ repeater server.

40. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 39, wherein the reflector comprises the cost determination mechanism.

41. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 39, wherein the reflector is co-located with an origin server from which the first resource originates.

42. (Currently amended) The ~~distributed hosting framework~~ content delivery system as recited in claim 39, wherein the cost comprises information relating a location associated with the particular client machine to each of the set of shared ~~content~~ repeater servers.

43-44. (canceled)

45. (Currently amended) A method, in a ~~distributed hosting framework~~ content delivery system operative in a computer network for delivering content to client machines, the computer network ~~and~~ comprising a plurality of origin servers, each of said origin servers having resources associated therewith~~, at least one name server that provides domain name service (DNS) resolution,~~ and the content delivery system comprising at least one shared ~~content~~ repeater server operable to replicate resources associated with the plurality of origin servers, the method comprising:

associating the at least one repeater server with a first alias name ~~at least some resources associated with a first origin server of plurality of origin servers with a first alias name that causes the at least one name server to indentify the at least one shared content repeater server to serve the resources associated with the first origin server in response to requests from client machines~~, wherein requests for a first resource located on a first origin server are directed, based at least in part on said first alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server; and

associating the at least one repeater server with a second alias name ~~at least some resources associated with a second origin server of said plurality of origin servers with a second alias name that causes the at least one name server to identify the at least one shared content repeater server to serve the resources associated with the second origin server in response to requests from client machines~~, wherein requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server.

46. (Currently amended) The method of claim 45, ~~wherein the associating step comprises~~ further comprising:

in response to a request for ~~a particular~~ the first resource issued by a particular client machine, identifying the at least one shared ~~content~~ repeater server from a set of shared ~~content~~ repeater servers based at least in part on a network cost of delivering resources to the particular client machine; and

associating the first resource with the first alias name.

47. (Currently amended) The method of claim 45, ~~wherein the associating step comprises~~ further comprising:

in response to a request for ~~a particular~~ the first resource issued by a particular client machine, identifying the at least one ~~content~~ repeater server from a set of shared ~~content~~ repeater servers based at least in part on load characteristics of the set of shared ~~content~~ repeater servers; and

associating the first resource with the first alias name.

48. (Currently amended) The method of claim 45, ~~wherein the associating step comprises~~ further comprising:

in response to a request for the first ~~a particular~~ resource issued by a particular client machine, identifying the at least one ~~content~~ repeater server from a set of shared ~~content~~ repeater servers based at least in part on a location of the particular client machine relative to the at least one shared ~~content~~ repeater server; and

associating the first resource with the first alias name.

49. (previously presented) The method as recited in claim 45, wherein the resources comprise an HTTP-based resource.

50. (previously presented) The method as recited in claim 45, wherein the resources comprise a File Transfer Protocol (FTP)-based resource.

51. (Canceled).

Please add the following new claims:

52. (NEW) A content delivery system as recited in claim 39 wherein the cost determination mechanism comprises a netmap determination means.

53. (NEW) A content delivery system as recited in claim 29, further comprising:

a subscriber table listing a plurality of content providers authorized to have resources delivered to client machines via the at least one repeater server, wherein the plurality of content providers comprises the first content provider and the second content provider; and

wherein the at least one repeater server is further constructed and adapted to analyze, using the subscriber table, information about an alias name received with

a client request for a resource to determine a content provider associated with the requested resource.

54. (New) The content delivery system as in claim 45, further comprising:

providing a subscriber table listing origin servers having content located thereon that is authorized for delivery to client machines via the at least one shared repeater server, wherein the origin servers comprise the first origin server and the second origin server; and

wherein the at least one repeater server is further constructed and adapted to analyze, using the subscriber table, information about an alias name received with a client request for a resource to determine an origin server associated with the requested resource.

55. (NEW) A content delivery system as recited in claim 29, wherein at least the first alias name is a domain name, the system further comprising:

at least one name server that provides domain name service (DNS) resolution on the first alias name to identify the at least one shared repeater server to deliver the first resource in response to a client request.

56. (NEW) A method as recited in claim 45, wherein at least the first alias name is a domain name and the computer network further comprises at least one name server that provides domain name service (DNS) resolution, the method further comprising:

receiving, at the at least one name server, the first alias name as part of a client request; and

resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to deliver the first resource in response to the client request.

57. (NEW) A method, in a content delivery system operative in a computer network for delivering content to client machines and comprising at least one shared repeater server operable to replicate resources stored on a plurality of origin servers, the method comprising:

associating a plurality of alias names with the at least one shared repeater server, each of said plurality of alias names being associated with an origin server, wherein a first alias name of said plurality of alias names is associated with a first origin server, and at least a second alias name of said plurality of alias names is associated with a second origin server distinct from said first origin server,

wherein requests for a first resource originating on the first origin server are directed, based at least in part on said first alias name, to the at least one shared repeater server for delivery of the first resource from said at least one repeater server; and

wherein requests for a second resource originating on the second origin server are directed, based at least in part on said second alias name, to the at least one shared repeater server for delivery of the second resource from said at least one repeater server.

REMARKS

Reconsideration and allowance of this application are respectfully requested.

By this Response, claims 27, 29-42, and 45-48 have been amended, claims 43, 44, and 51 have been canceled, and new claims 52-57 have been added.

No new matter has been added by these amendments.

Claims 27-42, 45-50, and 52-57 are pending in this application.

THE CLAIM REJECTIONS UNDER §101

The Examiner rejected claims 27-44 under 35 U.S.C. § 101 as being directed to non-statutory subject matter. Specifically, the Examiner stated that "Claims 27-44 are claimed as being a "distributed hosting framework " that is comprised only of software elements. As the claims are comprised only of software, they are treated as software per se. Software per se does not fit into any of the statutory categories of invention."

The grounds for this rejection are respectfully traversed.

The term "*per se*" means "Of, in, or by itself; standing alone, without reference to additional facts." (Blacks Law Dictionary, 2nd Pocket Edition, pg. 525).

First, the claims are not, as the Examiner would have it, directed to "software *per se*." Claims 27 and 28, *e.g.*, recite a distributed hosting framework that includes, amongst other things, "a set of repeater servers," "at least one first level name server," and "at least one second level name server." Claim 29 and its dependents recite a "content delivery **system**," that includes "at least one shared repeater server."

Second, even if, for the sake of argument, the claims are directed to "software, per se" the Examiner has provided no legal basis for his rejection. The Examiner has not shown or explained why **the claims** are not statutory. The

Examiner is required to provide the basis for his rejection of the claims, other than a cursory and conclusory statement.

In view of the above, withdrawal of this rejection under §101 is respectfully requested.

THE §112 REJECTIONS

The Examiner rejected claims 27-51 under 35 U.S.C. § 112, ¶2, as being indefinite.

According to the Examiner:

> Claim 27 recites the limitation "the embedded object identified by the modified embedded object URL" in the final limitation of the claim. There is insufficient antecedent basis for this limitation in the claim. Specifically, not [*sic*] embedded object appears to have been "identified" previously in the claim. ... Claims 29 and 45 feature the same problem as claim 27 and are rejected for the same reason as claim 27. The dependent claims are rejected as being dependent upon an indefinite claim."

Claim 27 has been amended to change the phrase "the embedded object" to "an embedded object". However, Applicant finds no similar claim language in claims 29 and 45 (claims 29 and 45 do not include the phrase "embedded object"). If the Examiner believes that there is indefinite language in those claims, he is respectfully requested to identify the specific language.

In view of the above, withdrawal of the rejection under §112 is respectfully requested.

THE PRIOR ART REJECTIONS

THE REJECTIONS UNDER §102

The Examiner rejected claims 27-38, 43-45, 47, and 49-51 under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 5,867,706 to Martin et al. (hereinafter "Martin"). The grounds for this rejection are respectfully traversed.

Martin teaches a system for delivering a single content provider's resources to clients from a so-called "parallel web server" (10) – a singled server having multiple processors (20, 30, 40). *See Martin* Fig. 1 (reproduced below) and col. 5, lines 10-12. Each of processors (20, 30, 40) in Martin's server 10 is assigned a different name. In the example described in Martin, processor 20 is named *www1.ibm.com*, processor 30 is named *www2.ibm.com*, and processor 40 is named *www3.ibm.com*. *Id.* at col. 5, lines 12-15. Note that in Martin the parallel web server 10 has only one name (e.g., *www.ibm.com*) and each processor 20, 30, 40 has only one name.

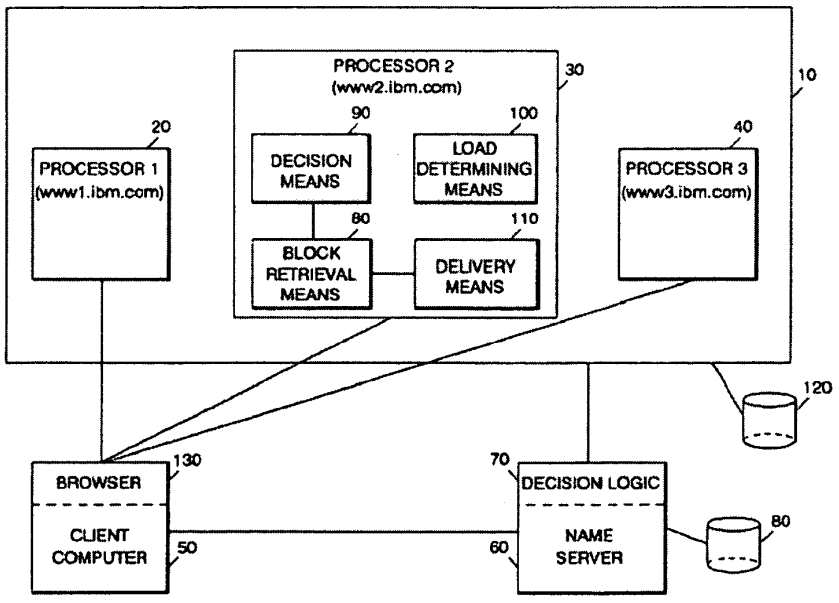


FIG. 1

The "object [of Martin is] ... to provide a mechanism to alleviate unbalanced load across the processors, thereby helping to avert ... [a] toasting effect." *Martin*, col. 2, lines 10-12. Thus, in Martin, requests made to the parallel web server (10) are distributed among the processors (20, 30, 40) using a processor load balancing technique that takes into account load on the various processors. See, e.g., *Id.* col. 5, lines 13-32. Martin describes the use of a name server to perform aspects of the processor load balancing. This approach is akin to that described in the "Background" section of the present application ("Once a Web site is served by multiple servers, a challenge is to ensure that the load is appropriately distributed or balanced among those servers. Domain name-server-based round-robin address resolution causes different clients to be directed to different mirrors." *Published Application No. 20050198334*, ¶0008; see also "... load balancing, takes into account the load at each server (measured in a variety of ways) to select which server should handle a particular request." *Id.* ¶0009).[1]

Martin describes a system that serves resources only from and on behalf of a single content source. In Martin, no matter which processor 20, 30 or 40 is used, Martin only teaches delivering resources from and on behalf of a single content provider – IBM in the example given. Martin's parallel server 10 has a single storage device 120 that is shared by the various processors 20, 30, 40. Thus, regardless of which processor is used for a particular client request, that processor will get the content from the same content source (that is, from storage device 120). This configuration is similar to a mirrored system in which each processor has an identical copy of the content, and it does not include features that allow its shared use by and on behalf of many content sources / providers at the same time. (Although in Martin no processor has its own copy of the content to be delivered.) In this sense, the teachings of Martin are akin to those described in the

[1] For the Examiner's convenience, citations to the present application made herein are to U.S. Publication No. 20050198334, the published version of the present application.

Background section of the present application. (See "A number of solutions have been developed to advance replication and mirroring. In general, these technologies are designed for use by a single Web site and **do not include features that allow their components to be shared by many Web sites simultaneously**." U.S. Publication No. 20050198334, ¶0006, emphasis added)

An additional teaching of Martin is that once a client request has been assigned to one of the processors (20, 30, 40) in the parallel web server 10, so-called "follow-on links" to the same server are also made via that previously-selected processor. As noted by Martin (*Martin*, col. 6, lines 17 to 26):

> ... the parallel Web server of the preferred embodiment is provided with the ability to modify the content of Web pages in the server as they are served, and takes advantage of that facility to modify the name of the processors to which the follow-on links refer. Considering the FIG. 1 example, since www1.ibm.com, www2.ibm.com and www3.ibm.com are all servers in the www.ibm.com parallel Web server, a follow-on link which would otherwise refer to www.ibm.com can be changed to refer to a specific processor, for example www1.ibm.com.

Martin gives an example of how the links to follow-on Web pages are modified, *e.g.*, at col. 7, lines 50-67. Specifically, Martin describes how (if a follow on Web page is called "second_page.html"), then the exemplary link:

http://www.ibm.com/second_page.html

will be changed to read

http://www1.ibm.com/second_page.html

As can be seen from the above, Martin changes the domain name of the link from "www.ibm.com" to "www1.ibm.com", but requires nor makes any other changes to the link.

The presently claimed invention is patentable over Martin for a number of reasons.

Independent Claim 29 recites a content delivery system having, *inter alia*, "at least one **shared** repeater server" that replicates at least some of the resources associated with a first content provider and a second content provider. (See also "for delivering resources associated with a plurality of content providers") This feature is not taught or in any way suggested by Martin. To the contrary, as noted above, Martin pertains to a server that serves resources only from and on behalf of a single content source (or content provider). Claim 29 explicitly recites a server infrastructure configured to serve resources to clients on behalf of multiple, distinct content sources (phrased in claim 29 as "content providers"). Similar arguments apply to **independent claim 45** (the content sources are referred to in claim 45 as "origin servers").

For at least these reasons, **claims 29 and 45** and their respective dependents are patentable over Martin. Furthermore, both claims explicitly require at least one repeater server to be associated with multiple alias names corresponding to either multiple, distinct content providers (as in claim 29) or multiple, distinct origin servers (as in claim 45). This feature – associating multiple alias names with a repeater server – supports the "shared server" aspect of the system of claims 29 and 45. The use of different alias names allows the system of the present invention to ascertain which content provider is associated with a request, and this is used by the system for a variety of purposes relevant to support a shared infrastructure, such as, among other things, providing the repeater servers of the presently claimed invention with information for use in obtaining requested resources that they do not already have, as well as creating logs concerning which

resources have been requested. Since Martin does not even suggest the applicability of its system in the delivery of resources on behalf of more than one distinct content source, Martin does not have any need for alias names, much less get into the specific alias-based approach recited in claims 29 and 45 ("first alias name ... [and a]... second alias name being associated with said at least one shared repeater server.").

The Examiner relies on Martin, col. 6, ll. 17-33 to support the allegation that Martin teaches the use of alias names (e.g., *www1.ibm.com, www2.ibm.com, www3.ibm.com*). However, as noted above, these names are not "alias" names, as recited in claims 29 and 45. Rather, in Martin, these are the actual names for the processors 20, 30 and 40 used by the same web server 10 to service client requests for resources, which web server 10 is exclusively taught in Martin as dedicated to a single content source. In other words, each of the three processors 20, 30, 40 is and can only be addressed by its single name (resolved by Martin's name server 60. *See*, Martin, col. 5, lines 37-42).

In summary, claims 29 and 45 recite an inventive approach that is certainly not anticipated by Martin and, given the contextual differences in the problems to be solved, respectfully, would not be rendered obvious in view of Martin, alone, modified or combined with any other proposed reference. Claims 30-44 and 46-56 depend from claims 29 and 45, respectively, and therefore are believed neither anticipated nor rendered obvious in view of Martin for at least the reasons so stated, as well as further reasons that are evident upon review of these claims.

Claim 30 depends from claims 29 and is therefore patentable over Martin for at least the reasons given above. Claim 30 recites a content delivery system as recited in claim 29, "wherein the least one shared repeater server is addressable using the first alias name and the second alias name." This feature is not taught or in any way suggested by Martin.

In Martin, the Web server 10 is *only* addressable using the name "*www.ibm.com*", and each of the three processors is *only* addressable by its given name (*www1.ibm.com* for processor 1, or *www2.ibm.com* for processor 2, or *www3.ibm.com* for processor 3). To the extent the Examiner considers *www1.ibm.com, www2.ibm.com*, and *www3.ibm.com* to be the aliases, then it necessarily follows that no processor in the parallel web processor is addressable using more than one of these names.

For at least this additional reason, claim 30 is further patentable over Martin, alone or in any proposed combination.

Claim 31 depends from claims 29 and is therefore patentable over Martin for at least the reasons given above. Claim 31 recites a content delivery system as recited in claim 29, including "a reflector constructed and adapted to associate the resources with the first alias name based, at least in part, on a relative cost of transmitting the resources to the client machines."

Martin describes a system that tries to perform processor load balancing, and all that Martin is concerned with is the load on the processors. So, in Martin, client requests are directed to one of the processors according to those loads (see, e.g., "decision means 90 refers to a load distribution record providing information about the loads on all the processors of the server." *Martin*, col. 7, lines 3 *et seq.*). But Martin has no teaching or suggestion of any processing "based, at least in part, on a relative cost of transmitting the resources to the client machines."

For at least these additional reasons, claim 31 is further patentable over Martin, alone or in any proposed combination.

Claim 46 depends from claims 45 and is therefore patentable over Martin for at least the reasons given above. Claim 46 recites a method as in claim 45, further comprising: "in response to a request for the first resource issued by a particular client machine, identifying the at least one shared repeater server from a

set of shared repeater servers based at least in part on a network cost of delivering resources to the particular client machine."

As noted above (for claim 31), Martin has no teaching or suggestion of any kind of selection "based ... on a network cost of delivering resources to the particular client machine." In Martin all selection is made based on processor loads.

For at least these additional reasons, claim 46 is further patentable over Martin, alone or in any proposed combination.

Independent Claim 27 recites a system (a distributed hosting framework) including "a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path." This modified URL is then used to identify a content server from which to serve the embedded object. In contrast, Martin creates so-called "follow-on links" "which would otherwise refer to www.ibm.com." *Martin*, col. 6, lines 17-33. As discussed above, Martin gives an example of how the links to follow-on Web pages are modified, *e.g.*, at col. 7, lines 50-67. Specifically, Martin describes how (if a follow on Web page is called "second_page.html"), then the exemplary link:

http://www.ibm.com/second_page.html URL#1

is changed to read

http://www1.ibm.com/second_page.html URL#2

In this example, in URL#1, the hostname is "www.ibm.com", and in URL#2, the hostname is "www1.ibm.com." As can be seen from the above, Martin changes the domain name of the link (URL#1) from "www.ibm.com" to

"www1.ibm.com", but requires nor makes any other changes to the link. In particular, Martin does not insert anything between the domain name (www.ibm.com) and the path ("/second_page.html"), and Martin does not **prepend** anything to "a domain name and path."

Even if the Examiner considers the "www" and "www1" in these URLs to be hostnames, then there is still no prepending in Martin of a hostname to a domain name and path.

The present application describes modifying (or rewriting) URLs, *e.g.*, http://<server>/<path> to http://<repeater>/<server>/<path>. (See, *e.g.*, ¶¶0072-73), where "<server>" refers to an origin server (e.g., a content provider's server), and "<repeater>" is a repeater server selected to handle a particular request. It should be noted that a reason for including the original "origin server" domain name in the modified URL in accordance with this embodiment of Applicant's invention is so that the repeater server knows where to go to get requested resources that it does not already have. While there are other ways to achieve this result, this feature is not required by a system such as Martin's in which all resources are obtained from the same place (*e.g.*, from storage 120). Since Martin's parallel web server is not shared, it has no need to determine where to go to get requested resources – it always goes to the same place, and that place does not vary for different requests to the server.

Martin thus teaches a system in which only a domain name, and <u>not additionally a host name</u>, is used to identify a processor 20, 30 or 40 for use in serving a resource to a requesting client.

Since Martin does not describe a shared system, he would have no need for any information "prepended to a domain name and path."

Additionally, in view of the lack of teaching in regards to both a domain name and a host name, Martin necessarily cannot teach a routine in which URLs are modified by the explicitly-recited prepending approach. Martin therefore as a

matter of law cannot anticipate claim 27, nor can it anticipate claim 28, which depends from and incorporates all limitations of claim 27.

Moreover, Martin fails to provide a motivation upon which to render claim 27 obvious. The use of dual names is a significant distinction between the embodiment of the invention recited in claim 27 and Martin's system. First and foremost, Martin is focused on load balancing amongst processors 20, 30 and 40 within a parallel web server 10 that serves resources on the World Wide Web. The web server 10 is, effectively, the origin server described throughout the present application. After a processor 20, 30 or 40 is selected by way of load balancing, the domain name of that processor is sent to the requesting client.

The invention recited in claim 27 directs client requests away from origin servers and to repeater servers. In Martin's system, a requested resource will always be served from web server 10 (albeit using different processors), whereas the invention recited in claim 27 results in the service of client requests away from originating web servers (e.g., 10) altogether and off-loads those processing tasks to a repeater server.

The invention recited in claim 27 alleviates the need for web servers such as the one taught by Martin. See U.S. Publication No. 20050198334, ¶0023 ("As a consequence of this rewriting process, when the client requests other resources based on identifiers in the particular requested resource, the client will make those requests directly to the selected repeater, **bypassing the reflector and origin server entirely**.") (emphasis added). Given this, its evident that the invention recited in the claims of the present application actually obviate the need to load balance Martin's web server in the first place (given an intent to bypass the origin server), thereby further proving that those of skill in the art in February, 1998 would not have been motivated to modify Martin in the direction of the claims recited in the present application.

In view of the above, withdrawal of the claim rejections under §102 are respectfully requested.

THE REJECTIONS UNDER §103

The Examiner rejected claims 39-42, 46, and 48 under 35 U.S.C. §103(a) as being unpatentable over Martin in view of U.S. Patent Number 6,052,718 to Gifford et al. (hereinafter "Gifford"). The grounds for this rejection are respectfully traversed.

Claims 38-42, 46 and 48 each depend from either independent claim 29 or independent claim 45, and therefore, incorporate at least those limitations described above as being deficient from Martin. The combination of Martin and Gifford therefore fails to constitute a *prima facie* case of obviousness and respectfully should be withdrawn. In addition, the problems to be solved by Martin and the context to which Martin relates fails altogether to provide any motivation or suggestion to modify or improve Martin in the direction of any invention claimed in the present application.

The Examiner acknowledges that "Martin does not teach choosing a server based on the performance for a particular client." However, according to the Examiner, "Gifford teaches the selection of a load balanced server based on the performance that it can provide a particular client system (See Figures 4A and 4B and corresponding text). It would have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Martin regarding a method for load balancing embedded objects with the teachings of Gifford regarding selecting a load balanced server that performs well for a particular client because doing so enhances performance for clients (col. 1, lines 42-57 of Gifford)."

Gifford uses so-called replica routers and a replica routing database to route clients requests to replica servers. (See, *e.g.*, Gifford, Fig. 4A and related description). But Gifford has no teaching or notion of the shared repeater servers

of the presently claimed invention, and thus does not overcome the deficiencies in Martin.

In view of the above, withdrawal of the rejections under §103 are respectfully requested.

NEW CLAIMS 52-57

New claims 52-56 each depend from either independent claim 29 or independent claim 45, and therefore, incorporate at least those limitations described above as being deficient from Martin.

In addition, **Claims 53 and 54** further distinguish the shared infrastructure context void from Martin, as described above in greater detail.

Firstly, Martin fails altogether to teach the use of a table for any reason or any purpose, much less a table that lists a plurality of content providers or origin servers, as the case may be. *Secondly*, because Martin fails altogether to teach an infrastructure that is shared by distinct entities, Martin has no need for and provides absolutely no motivation for the use of such a table.

Claim 52 recites the Netmap procedure detailed in the specification (see, e.g., ¶¶0152 et seq.) used to implement the cost determination mechanism. Indeed, claim 52 recites "netmap determination means" to reference this procedure by invoking 35 U.S.C. §112, paragraph 6. It should be appreciated that the Applicant has not invoked 35 U.S.C. §112, paragraph 6 with respect to any other term in any claim of the present application.

Claims 55 and 56 further define at least the first alias name as a domain name and recite the use of a specific device (a domain name server) commonly known in the art for resolving domain names (e.g., U.S. Publication No. 20050198334, ¶0051 and ¶0055) to resolve the first alias name to identify the at least one shared repeater server. It is brought to the Examiner's attention that the

limitations of claims 55 and 56 were previously recited in independent claims 29 and 45, but are hereby moved into dependent form so as to not limit independent claims 29 and 45 to this specific domain name service implementation.

Independent claim 57 recites "associating a plurality of alias names with the at least one shared repeater server, each of said plurality of alias names being associated with an origin server, wherein a first alias name of said plurality of alias names is associated with a first origin server, and at least a second alias name of said plurality of alias names is associated with a second origin server distinct from said first origin server." The invention of claim 57 is patentable over Martin (alone or in any proposed combination) for at least the following reasons:

- Martin does not associate a plurality of alias names with any server (or processor);
- Martin does not teach or in any way suggest a shared repeater server; and
- Martin has no teaching or suggestion of a first alias name being associated with a first origin server and a second alias name being associated with a second origin server distinct from the first origin server.

Accordingly, claim 57 is believed allowable over Martin.

OTHER MATTERS

Applicant does not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application

The Examiner is respectfully advised that a Court has construed certain terms in litigation involving a parent of the present application. The Markman

Order from that case is cited in the Information Disclosure Statement being filed herewith.

CONCLUSION

Applicant respectfully submits that this application is in condition for allowance. An early action to that effect is earnestly solicited.

The Examiner is kindly requested to contact the undersigned at the number provided to schedule a personal interview to resolve any outstanding issues in this case.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 75948	Respectfully submitted, By: **/Brian Siritzky/Reg. No. 37,497** Brian Siritzky, Ph.D. Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

PTO/SB/06 (07-06)
Approved for use through 1/31/2007. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/065,412	02/23/2005	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☐		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)	SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
	07/09/2009	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	* 29	Minus	** 48	= 0	X $ =		OR	X $52=	0
	Independent (37 CFR 1.16(h))	* 4	Minus	***5	= 0	X $ =		OR	X $220=	0
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)					
		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =		OR	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =		OR	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".

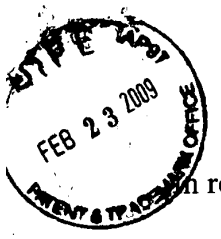
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

Legal Instrument Examiner:
/TYWANA P. LOVELACE/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.





IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:		Atty. Docket No.:	2711-0073
	FARBER, David et al.	Confirmation No.:	3296
Appln. No.:	11/065,412	Group Art Unit:	2442
Filed:	February 23, 2005	Examiner:	BLAIR, Douglas
Title:	SHARED CONTENT DELIVERY INFRASTRUCTURE	Date:	February 23, 2009

REPLY/AMENDMENT/LETTER TRANSMITTAL COVER SHEET

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a reply/amendment/letter in the above-identified application and includes the attachments hereto. The signature below is treated as the signature to the attachments in the absence of a signature thereto.

FEE REQUIREMENTS

	Claims remaining after amendment		*Highest number previously Paid For*		*Present Extra*	*Rate: Large/Small Entity*	
1. *Total*		minus		*	=	X $50 / $ 25 =	$
			* If this number is less than 20, enter "20"				
2. *Independent Claims*		minus		**	=	X $210 / $105 =	$
			** If this number is less than 3, enter "3"				
3. If amendment enters multiple dependent claim(s) into this application for first time (leave this line blank if this is an reissue application)						$370 / $185 =	$
4. **Original due date:**							
				Check time period and enter appropriate fee			
5. **Petition is hereby made** to extend the original due date to cover the date this response is filed for which the requisite fee is:				☐ 1 month ☐ 2 months. ☐ 3 months ☐ 4 months ☐ 5 months.		$120 / $60 $460 / $230 $1050 / $525 $ 1640/ $820 $2230 / $1115	$
6. ☐ Attached is a Petition/Fee under Rule No.						$	$
7. Other Fee for : Information Disclosure Statement under 37 C.F.R. § 1.17(p)							$180
8. **Total Fee Enclosed:**							$180

This is a 2 page form, continued on next page . . .

9. ☒ Please charge the total fee on line 8 above to our deposit account below under the stated order number.

CHARGE STATEMENT: Deposit Account No. 501860, order no. (client-matter no.) **2711-0073.**
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/ Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached. **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 75948	Respectfully submitted, By: _____ Brian Siritzky, Ph.D. Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: February 23, 2009

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

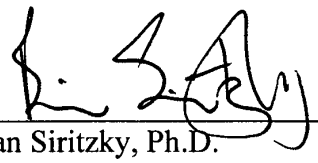
Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference listed on the attached PTO-1449. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This IDS is intended to be in full compliance with the rules, but should the Examiner find any part of its requirement content to have been omitted, prompt notice to that effect is earnestly solicited, along with additional time under Rule 97(f), to enable Applicant to comply fully.

02/24/2009 DEMMANU1 00000046 501860 11065412
01 FC:1806 180.00 DA

This Information Disclosure Statement is being filed more than three (3) months after the U.S. filing date AND after the mailing date of the first Office Action on the merits, but before the mailing date of a Final Rejection or Notice of Allowance. Please credit or debit Deposit Account No. 501860 the amount required under 37 C.F.R. § 1.17(p) as needed to ensure consideration of the disclosed information.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 • FAX: (703) 894-6430	



INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 1	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	1-1	US-5777988	Jul-98	Cisneros
	1-2	US-5940831	08-1999	Takano
	1-3	US-6324580	11-2001	Jindal
	1-4	US-6351775	02-2002	Yu
	1-5			
	1-6			
	1-7			
	1-8			
	1-9			
	1-10			
	1-11			
	1-12			
	1-13			
	1-14			
	1-15			
	1-16			
	1-17			
	1-18			
	1-19			
	1-20			
	1-21			
	1-22			
	1-23			
	1-24			
	1-25			
	1-26			

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/065,412	02/23/2005	David A. Farber	2711-0073	3296

75948 7590 01/09/2009

DAVIDSON BERQUIST JACKSON & GOWDEY, LLP
ATTN: BRIAN SIRITZKY, Ph.D.
4300 WILSON BLVD., 7TH FLOOR
ARLINGTON, VA 22203

EXAMINER	
BLAIR, DOUGLAS B	
ART UNIT	PAPER NUMBER
2442	

MAIL DATE	DELIVERY MODE
01/09/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	11/065,412	FARBER ET AL.	
	Examiner	Art Unit	
	DOUGLAS B. BLAIR	2442	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *10 December 2008*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *27-51* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *27-51* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *9/23/08,12/22/08,12/30/08*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

> Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

Claims 27-44 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. Claims 27-44 are claimed as being a "distributed hosting framework" that is comprised only of software elements. As the claims are comprised only of software, they are treated as software per se. Software per se does not fit into any of the statutory categories of invention.

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 27-51 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claim 27 recites the limitation "the embedded object identified by the modified embedded object URL" in the final limitation of the claim. There is insufficient antecedent basis for this limitation in the claim. Specifically, not embedded object appears to have been "identified" previously in the claim.

Claims 29 and 45 feature the same problem as claim 27 and are rejected for the same reason as claim 27. The dependent claims are rejected as being dependent upon an indefinite claim.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 27-38, 43-45, 47, 49-51 are rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Patent Number 5,867,706, to Martin et al.

As to claim 27, Martin teaches a distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising: a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path **(col. 6, lines 17-33)**; a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server **(col. 6, lines 17-33)**; at least one first level name server that provides a first level domain name service (DNS) resolution **(col. 5, line 61-col. 6, line 16)**; and at least one second level name server that provides a second level domain name service (DNS) resolution **(col. 5, line 61-col. 6, line 16)**; wherein in response to requests for the web page, generated by the client machines the web page including the modified

embedded object URL is served from the content provider server **(col. 6, lines 17-33)** and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers **(col. 5, line 61-col. 6, line 16)**.

As to claim 28, Martin teaches the distributed hosting framework as recited in claim 27, further including a redundant second level name server **(col. 5, line 61-col. 6, line 16)**.

As to claim 29, Martin teaches a distributed hosting framework operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, the framework comprising: at least one shared content server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers **(col. 6, lines 17-33)**; and at least one name server that provides domain name service (DNS) resolution **(col. 5, line 61-col. 6, line 16)**, wherein resources associated with the first content provider are associated with a first alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the first content provider made from said client machines **(col. 6, lines 17-33)**; and wherein the resources associated with the second content provider are associated with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the second content provider made from said client machines **(col. 6, lines 17-33)**.

As to claim 30, Martin teaches the distributed hosting framework as recited in claim 29, wherein the least one shared content server is addressable using the first alias name and the second alias name **(col. 6, lines 17-33)**.

As to claim 31, Martin teaches the distributed hosting framework as recited in claim 29, further comprising: a reflector constructed and adapted to associate the resources with the first alias name based, at least in part, on a relative cost of transmitting the resources to the client machines **(col. 6, line 56-col. 9, line 9)**.

As to claim 32, Martin teaches the distributed hosting framework as recited in claim 29 further comprising: a set of shared content servers that comprises the at least one content server **(col. 6, lines 17-33)**; and a reflector constructed and adapted to associate a particular resource with the first alias name after selecting the at least one shared content server from the set of shared content servers in response to a request for the particular resource issued by a particular client machine **(col. 6, lines 17-33)**.

As to claim 33, Martin teaches the distributed hosting framework as recited in claim 29, wherein at least one of the resources is an embedded object **(col. 6, lines 17-33)**.

As to claim 34, Martin teaches the distributed hosting framework as recited in claim 29, wherein the requests comprise host tag names that identify specific origin servers from which the resources originate **(col. 6, lines 17-33)**.

As to claim 35, Martin teaches the distributed hosting framework as recited in claim 34, wherein the host tag names are included in HTTP headers **(col. 6, lines 17-33)**.

As to claim 36, Martin teaches the distributed hosting framework as recited in claim 32, further comprising: a load determination mechanism constructed and adapted to characterize load on the set of content servers for use by the reflector in selecting the at least one shared content server **(col. 6, line 56-col. 9, line 9)**.

As to claim 37, Martin teaches the distributed hosting framework as recited in claim 36, wherein the reflector comprises the load mechanism **(col. 6, line 56-col. 9, line 9)**.

As to claim 38, Martin teaches the distributed hosting framework as recited in claim 37, wherein the reflector is co-located with an origin server **(col. 6, line 56-col. 9, line 9)**.

As to claim 43, Martin teaches the distributed hosting framework as recited in claim 29, further including a redundant name server **(col. 5, line 61-col. 6, line 16)**.

As to claim 44, Martin teaches the distributed hosting framework as recited in claim 29, wherein the at least one name server comprises: at least one first level name server that provides a first level domain name service (DNS) resolution; and at least one second level name server that provides a second level DNS resolution **(col. 5, line 61-col. 6, line 16)**.

As to claim 45, Martin teaches a method, in a distributed hosting framework operative in a computer network for delivering content to client machines and comprising a plurality of origin servers, each of said origin servers having resources associated therewith, at least one name server that provides domain name service (DNS) resolution **(col. 5, line 61-col. 6, line 16)**, and at least one shared content server operable to replicate resources associated with the plurality of origin servers, the method comprising: associating at least some resources associated with a first origin server of plurality of origin servers with a first alias name that causes the at least one name

server to identify the at least one shared content server to serve the resources associated with the first origin server in response to requests from client machines **(col. 6, lines 17-33)**; and associating at least some resources associated with a second origin server of said plurality of origin servers with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources associated with the second origin server in response to requests from client machines **(col. 6, lines 17-33)**.

As to claim 47, Martin teaches the method of claim 45, wherein the associating step comprises: in response to a request for a particular resource issued by a particular client machine, identifying the at least one content server from a set of shared content servers based at least in part on load characteristics of the set of shared content servers **(col. 6, line 56-col. 9, line 9)**.

As to claim 49, Martin teaches the method as recited in claim 45, wherein the resources comprise an HTTP-based resource **(col. 6, lines 17-33)**.

As to claim 50, Martin teaches the method as recited in claim 45, wherein the resources comprise a File Transfer Protocol (FTP)-based resource **(col. 1, lines 25-30)**.

As to claim 51, Martin teaches the method of claim 45 wherein the at least one name server comprises: at least one first level name server that provides a first level domain name service (DNS) resolution; and at least one second level name server that provides a second level DNS resolution **(col. 5, line 61-col. 6, line 16)**.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 39-42, 46, and 48 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent Number 5,867,706 to Martin et al. in view of U.S. Patent Number 6,052,718 to Gifford et al.

As to claims 39-42, 46, and 48, Martin teaches the claimed scenarios for load balancing objects however Martin does not teach choosing a server based on the performance for a particular client. Gifford teaches the selection of a load balanced server based on the performance that it can provide a particular client system (See Figures 4A and 4B and corresponding text). It would have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Martin regarding a method for load balancing embedded objects with the teachings of Gifford regarding selecting a load balanced server that performs well for a particular client because doing so enhances performance for clients (col. 1, lines 42-57 of Gifford).

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to DOUGLAS B. BLAIR whose telephone number is (571)272-3893. The examiner can normally be reached on 9:00am-5:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Andrew Caldwell can be reached on (571) 272-3868. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Douglas B Blair/
Primary Examiner, Art Unit 2442

.

Notice of References Cited	Application/Control No. 11/065,412	Applicant(s)/Patent Under Reexamination FARBER ET AL.	
	Examiner DOUGLAS B. BLAIR	Art Unit 2442	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,230,196	05-2001	Guenthner et al.	709/223
*	B	US-6,351,775	02-2002	Yu, Philip Shi-Lung	709/238
*	C	US-6,078,943	06-2000	Yu, Philip Shi-Lung	718/105
*	D	US-6,012,090	01-2000	Chung et al.	709/219
*	E	US-5,940,831	08-1999	Takano, Hajime	707/10
*	F	US-5,751,961	05-1998	Smyk, Darek Andrew	709/217
*	G	US-6,324,580	11-2001	Jindal et al.	709/228
*	H	US-5,867,706	02-1999	Martin et al.	718/105
*	I	US-6,052,718	04-2000	Gifford, David K.	709/219
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11065412	FARBER ET AL.
	Examiner	Art Unit
	DOUGLAS B BLAIR	2442

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	09/15/2008	01/02/2009							
	1		-							
	2		-							
	3		-							
	4		-							
	5		-							
	6		-							
	7		-							
	8		-							
	9		-							
	10		-							
	11		-							
	12		-							
	13		-							
	14		-							
	15		-							
	16		-							
	17		-							
	18		-							
	19		-							
	20		-							
	21		-							
	22		-							
	23		-							
	24		-							
	25	✓	-							
	26		-							
	27		✓							
	28		✓							
	29		✓							
	30		✓							
	31		✓							
	32		✓							
	33		✓							
	34		✓							
	35		✓							
	36		✓							

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11065412	FARBER ET AL.
	Examiner	Art Unit
	DOUGLAS B BLAIR	2442

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	09/15/2008	01/02/2009							
	37		✓							
	38		✓							
	39		✓							
	40		✓							
	41		✓							
	42		✓							
	43		✓							
	44		✓							
	45		✓							
	46		✓							
	47		✓							
	48		✓							
	49		✓							
	50		✓							
	51		✓							

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11065412	FARBER ET AL.
	Examiner	Art Unit
	DOUGLAS B BLAIR	2442

SEARCHED			
Class	Subclass	Date	Examiner
709	217,219,245	9/15/2008	DBB
	update	12/30/08	DBB

SEARCH NOTES		
Search Notes	Date	Examiner
East text search	9/15/2008	DBB
update	12/30/08	DBB

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	6	(choos$3 or pick$3 or select $4) with (cluster or farm) with (proximity or locat$3) near2 client$1	US-PGPUB; USPAT	OR	OFF	2009/01/04 18:28
L2	0	L1 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/01/04 18:28
L3	156	(choos$3 or pick$3 or select $4) near2 (server$1) with (proximity or locat$3) near2 client$1	US-PGPUB; USPAT	OR	OFF	2009/01/04 18:29
L4	15	L3 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/01/04 18:29
L5	0	L4 and (embed$4 near object $1)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2009/01/04 18:30
S1	0	((embed$4 near URL) or plugin) same prepend$3 near3 domain	US-PGPUB; USPAT	OR	OFF	2008/09/15 05:42
S2	77	prepend$3 near3 domain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:43
S3	3	S2 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:45
S4	14	prepend$3 with domain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46
S5	0	S4 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46

S6	0	((embed$4 near URL) or plugin) same prepend$3 near3 domain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46
S7	42	((embed$4 near URL) or plugin) same (domain with URL)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:47
S8	1	S7 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:47
S9	72	subdomain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:49
S10	5	S9 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:50
S11	0	((embed$4 near URL) or plugin) same subdomain$1	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
S12	18	"embedded object URL"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
S13	1	S12 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
S14	411	("20040139097" | "20050038851" | "20050114296" | "4591983" | "4594704" | "4726017" | "4803641" | "4839798" | "4847784" | "4920432" | "4922417" | "4943932" | "4949187" | "4949248" | "0449570" | "5029232" | "5130792" | "5132992" | "5136716" | "5172413" | "5191573" | "5253275" | "5253341" | "5287499" | "5287537" | "5291554" |	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:00

		"5341477"	"5371532"				
		"5410343"	"5414455"				
		"5442389"	"5442390"				
		"5442749"	"5471622"				
		"5475615"	"5508732"				
		"5515511"	"5519435"				
		"5528281"	"5539621"				
		"5542087"	"5544313"				
		"5544327"	"5550577"				
		"5550863"	"5550982"				
		"5557317"	"5572643"				
		"5590288"	"5592611"				
		"5594910"	"5603026"				
		"5619648"	"5623656"				
		"5625781"	"5627829"				
		"5630067"	"5633999"				
		"5634006"	"5638443"				
		"5644714"	"5646676"				
		"5649186"	"5659729"				
		"5666362"	"5671279"				
		"5675734"	"5682512"				
		"5699513"	"5712979"				
		"5715453"	"5721914"				
		"5734831"	"5740423"				
		"5742762"	"5751961"				
		"5761507"	"5761663"				
		"5764906"	"5774660"				
		"5774668"	"5777989"				
		"5784058"	"5796952"				
		"5799141"	"5802106"				
		"5802291"	"5812769"				
		"5815664"	"5826031"				
		"5828847"	"5832506"				
		"5832514"	"5835718"				
		"5838906"	"5845303"				
		"5856974"	"5862339"				
		"5867706"	"5867799"				
		"5870546"	"5870559"				
		"5878212"	"5884038"				
		"5890171"	"5893116"				
		"5894554"	"5896533"				
		"5903723"	"5907704"				
		"5913028"	"5913033"				
		"5918010"	"5919247"				
		"5920701"	"5933832"				
		"5935207"	"5945989"				
		"5956489"	"5956716"				
		"5959008"	"5961596"				
		"5966440"	"5968121"				
		"5978791"	"5983214"				
		"5983227"	"5987606"				
		"5991809"	"6002720"				
		"6003030"	"6006264"				
		"6012090"	"6014686"				
		"6014698"	"6018516"				
		"6026440"	"6029175"				
		"6029176"	"6035332"				
		"6038216"	"6038310"				

		"6038610" | "6041324" | "6044405" | "6046980" | "6049831" | "6052718" | "6052730" | "6065051" | "6065062" | "6070191" | "6081829" | "6092112" | "6092204" | "6105028" | "6108673" | "6108703" | "6112231" | "6112240" | "6115357" | "6119143" | "6125388" | "6128601" | "6128660" | "6134583" | "6144375" | "6144702" | "6151624" | "6154753" | "6163779" | "6167427" | "6173311" | "6173322" | "6178160" | "6181867" | "6185619" | "6189030" | "6226642" | "6230196" | "6243760" | "6256675" | "6266699" | "6269394" | "6282569" | "6286045" | "6298041" | "6314565" | "6347085" | "6360256" | "6370571" | "6370580" | "6412000" | "6415280" | "6421726" | "6456630" | "6480893" | "6484204" | "6496856" | "6502125" | "6502215" | "6557054" | "6581090" | "6591299" | "6611862" | "6658464" | "6665706" | "6741563" | "6799221" | "6915329" | "6928442" | "7047300").PN.				
S15	212	("20040139097" | "20050038851" | "20050114296" | "4591983" | "4594704" | "4726017" | "4803641" | "4839798" | "4847784" | "4920432" | "4922417" | "4943932" | "4949187" | "4949248" | "0449570" | "5029232" | "5130792" | "5132992" | "5136716" | "5172413" | "5191573" | "5253275" | "5253341" | "5287499" | "5287537" | "5291554" | "5341477" | "5371532" | "5410343" | "5414455" | "5442389" | "5442390" | "5442749" | "5471622" | "5475615" | "5508732" | "5515511" | "5519435" | "5528281" | "5539621" | "5542087" | "5544313" |	US-PGPUB; USPAT	OR	OFF	2008/09/15 06:00

		"5544327"	"5550577"				
		"5550863"	"5550982"				
		"5557317"	"5572643"				
		"5590288"	"5592611"				
		"5594910"	"5603026"				
		"5619648"	"5623656"				
		"5625781"	"5627829"				
		"5630067"	"5633999"				
		"5634006"	"5638443"				
		"5644714"	"5646676"				
		"5649186"	"5659729"				
		"5666362"	"5671279"				
		"5675734"	"5682512"				
		"5699513"	"5712979"				
		"5715453"	"5721914"				
		"5734831"	"5740423"				
		"5742762"	"5751961"				
		"5761507"	"5761663"				
		"5764906"	"5774660"				
		"5774668"	"5777989"				
		"5784058"	"5796952"				
		"5799141"	"5802106"				
		"5802291"	"5812769"				
		"5815664"	"5826031"				
		"5828847"	"5832506"				
		"5832514"	"5835718"				
		"5838906"	"5845303"				
		"5856974"	"5862339"				
		"5867706"	"5867799"				
		"5870546"	"5870559"				
		"5878212"	"5884038"				
		"5890171"	"5893116"				
		"5894554"	"5896533"				
		"5903723"	"5907704"				
		"5913028"	"5913033"				
		"5918010"	"5919247"				
		"5920701"	"5933832"				
		"5935207"	"5945989"				
		"5956489"	"5956716"				
		"5959008"	"5961596"				
		"5966440"	"5968121"				
		"5978791"	"5983214"				
		"5983227"	"5987606"				
		"5991809"	"6002720"				
		"6003030"	"6006264"				
		"6012090"	"6014686"				
		"6014698"	"6018516"				
		"6026440"	"6029175"				
		"6029176"	"6035332"				
		"6038216"	"6038310"				
		"6038610"	"6041324"				
		"6044405"	"6046980"				
		"6049831"	"6052718"				
		"6052730"	"6065051"				
		"6065062"	"6070191"				
		"6081829"	"6092112"				
		"6092204"	"6105028"				
		"6108673"	"6108703"				

		"6112231" | "6112240" | "6115357" | "6119143" | "6125388" | "6128601" | "6128660" | "6134583" | "6144375" | "6144702" | "6151624" | "6154753" | "6163779" | "6167427" | "6173311" | "6173322" | "6178160" | "6181867" | "6185619" | "6189030" | "6226642" | "6230196" | "6243760" | "6256675" | "6266699" | "6269394" | "6282569" | "6286045" | "6298041" | "6314565" | "6347085" | "6360256" | "6370571" | "6370580" | "6412000" | "6415280" | "6421726" | "6456630" | "6480893" | "6484204" | "6496856" | "6502125" | "6502215" | "6557054" | "6581090" | "6591299" | "6611862" | "6658464" | "6665706" | "6741563" | "6799221" | "6915329" | "6928442" | "7047300").PN.				
S16	197	S15 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:01
S17	12535	709/217,219,245.ccls.	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:07
S18	27	(redundant near dns)	US-PGPUB; USPAT	OR	OFF	2008/12/30 22:39
S19	24	S18 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:41
S20	34	(redundan$2 near2 (dns or "domain name service"))	US-PGPUB; USPAT	OR	OFF	2008/12/30 22:43
S21	24	S20 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:43

S22	79	prepend$3 near3 domain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S23	3	S22 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S24	14	prepend$3 with domain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S25	0	S24 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S26	42	((embed$4 near URL) or plugin) same (domain with URL)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S27	1	S26 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S28	77	subdomain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S29	5	S28 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S30	18	"embedded object URL"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S31	1	S30 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45

S32	0	(S23 S25 S27 S29 S31) and @pd>"20080801"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:45
S33	83	((redundan$2 or failover or backup) near2 (dns or "domain name service"))	US-PGPUB; USPAT	OR	OFF	2008/12/30 22:46
S34	24	S33 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:46
S35	6	("7054935" "6654807" "6473405").pn.	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 22:58
S36	2	"20050198334".pn.	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:01
S37	1	replicat$3 near "domain name"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:04
S38	3186	replicat$3 with (domain or URL or website)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:04
S39	878	S38 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:05
S40	5	S38 same subdomain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:05
S41	1	S40 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:06
S42	8482	(mirror$3 or cluster$1 or replicat$3) with (domain or URL or website)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:08

S43	13	S42 same subdomain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:08
S44	3	S43 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:08
S45	12	S42 same hostname	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:09
S46	0	S45 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:09
S47	312	S42 same (domain near (name or format$4))	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:10
S48	31	S47 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:10
S49	57	blair.xa. and douglas.xa.	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/12/30 23:36
S50	0	("7281032").URPN.	USPAT	OR	OFF	2008/12/30 23:37
S51	46	("4802164" | "4837680" | "5155845" | "5680574" | "5774660" | "5781908" | "5841980" | "5867706" | "5878218" | "5890156" | "5917998" | "5920725" | "5950203" | "6003045" | "6006264" | "6052718" | "6061685" | "6061761" | "6088330" | "6088766" | "6092213" | "6189071" | "6230190" | "6237063" | "6282671" | "6314458" | "6324580" | "6324692" | "6397292" | "6408369" | "6487634" | "6502205" | "6516350" | "6542962" |	US-PGPUB; USPAT; USOCR	OR	OFF	2008/12/30 23:37

		"6601101" | "6606643" | "6633538" | "6654752" | "6694406" | "6799245" | "6826613" | "6912668" | "6922724" | "7051066" | "7072979" | "7188157").PN. OR ("7281032").URPN.				
S52	1	"6185598".pn.	US-PGPUB; USPAT	OR	OFF	2009/01/04 17:24

1/4/09 6:49:19 PM
C:\Documents and Settings\dblair\My Documents\EAST\workspaces\11065412.wsp

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	1-1	Defendant Limelight Networks, Inc.'s Memorandum in Support of its Motion for Summary Judgment of Non-Infringement and Invalidity of all Patents-In-Suit, for Civil Action No. 2:07cv589 (RGD-FBS), August 28, 2008.	
/DB/	1-2	Interlocutory Decision in Opposition proceedings for EPO Appln. No. 00 128 346.4 - 2211, Patent No. 1143337, September 22, 2008.	
/DB/	1-3	Level 3 Communications, LLC's Non-Confidential Version of its Opposition Brief to Defendant's Motion for Summary Judgment of Non-Infringement and Invalidity of all Patents-In-Suit in Civil Action No. 2:07cv589 (JBF-FBS), Sept. 17, 2008.	
/DB/	1-4	Level 3's Memorandum of Law in Reply to Limelight's Claim Construction Brief for U.S. Patent Nos. 6,654,807; 7,054,935; and 6,473,405, in Civil File No. 2:07CV589, June 27, 2008.	
/DB/	1-5	Level 3's Memorandum of Law in Reply to Limelight's Claim Construction Brief, in Civil File No. 2:07CV589, July 3, 2008.	
/DB/	1-6	Memorandum of Law in Support of Level 3's Proposed Claim Constructions, from Case No. 2:07cv589 RDG FBS, June 13, 2008.	
/DB/	1-7	Minutes of Oral Proceedings Appln. No. 00 128 346.4 - 2211, Patent No. 1143337, September 22, 2008.	

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

OIPE IAP81 DEC 22 2008 PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 5	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	1-1	Amended Answer of Defendant Limelight Networks, Inc., Exhibit 1 to Defendant Limelight Network Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, in Civil Action No. 2:07cv589, September 9, 2008 [18 pgs].	
/DB/	1-2	Claim construction Opinion and Order in Civil Action No. 2:07:cv589, December 10, 2008 [51 pages].	
/DB/	1-3	Communications from European Patent Office in European Patent Application No. 99 906 680.6: Minutes of Oral Hearing of 11/20/2008, including auxiliary requests; and Invitation under Art. 94(3) and Rule 71(1) EPC (dated Dec. 11, 2008) to amended dependent claims. [32 pgs.]	
/DB/	1-4	Declaration Kevin C. Almeroth in Support of Limelight Networks, Inc.'s Motion for Summary Judgment of Non-Infringement and Invalidity of All Patents-In-Suit, filed August 28, 2008 in Civil Action No. 2:07cv589 (RGD-FBS) [39 pgs.].	
/DB/	1-5	Defendant Limelight Network Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, in Civil Action No. 2:07cv589, September 9, 2008 [4 pgs].	
/DB/	1-6	Defendant Limelight Network Inc.'s Memorandum Of Points And Authorities In Support Of Its Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses (PUBLIC VERSION), in Civil Action No. 2:07cv589, September 9, 2008 [18 pgs.]	
/DB/	1-7	Defendant Limelight Network Inc.'s Memorandum of Points and Authorities in Support of its Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, (Public Version, redacted) filed 09/09/208, in civil action No. 2:07-cv-00589-RGD-FBS.	

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	2-1	Defendant Limelight Network Inc.'s Rebuttal Brief In Support Of Its Combined Motion For Leave To Amend Its Answer And Compel Related Interrogatory Responses (Public version), filed 10/07/2008 in civil action No. 2:07-cv-00589-RGD -FBS.	
/DB/	2-2	Defendant Limelight Network Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (Redacted), April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [46 pgs.]	
/DB/	2-3	Defendant Limelight Networks, Inc.'s Claim Construction Brief for U.S. Patent Nos. 6,654,807; 7,054,935; and 6,473,405, filed in Civil Action No. 2:07CV589 (RGD-FBS), June 27, 2008 [38 pgs.].	
/DB/	2-4	Defendant Limelight Networks, Inc.'s Reply Brief In Support Of Its Motion For Summary Judgment Of Non-Infringement And Invalidity Of All Patents-In-Suit, filed October 1, 2008 in Civil Action No. 2:07cv589 (JBF-FBS) [21 pgs.].	
/DB/	2-5	Defendant Limelight Networks, Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (No. 2-4, 12-13, And 19) [Redacted], May 16, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [70 pgs.].	
/DB/	2-6	Defendant Limelight Networks, Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (Nos. 1-11, 13-14, 17 and 19-20, [redacted] 02/09/2008, in civil action No. 2:07-cv-00589-RGD-FBS [101 pgs.].	
/DB/	2-7	Exhibit A to Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [68 pgs.].	

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	3-1	Exhibit B to Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [20 pgs.].	
/DB/	3-2	Exhibits A-G of Plaintiff Level 3 Communications, LLC's Memorandum Of Law In Opposition To Defendant Limelight Networks, Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, 09/30/08, Civil Action No. 2:07cv589, (Exhibit A [7 pgs.], B [5 pgs.], C [47 pgs.], D [45 pgs.], E [7 pgs.], F [8 pgs.], G [5 pgs.]).	
/DB/	3-3	Exhibits A-V filed in support of Limelight Networks, Inc.'s Motion for Summary Judgment of Non-Infringement and Invalidity of All Patents-In-Suit, August 28, 2008 in Civil Action No. 2:07cv589 (JBF-FBS) [277 pgs.]	
/DB/	3-4	Level 3 Communications Supplemental Exhibit A in Civil Action No. 2:07cv589 (JBF -FBS) [69 pgs.]	
/DB/	3-5	Level 3 Communications, LLC's Confidential Version Of Its Opposition Brief To Defendant's Motion For Summary Judgment Of Non-Infringement And Invalidity Of All Patents-In-Suit, September 17, 2008, filed in Civil Action No. 2:07cv589 (JBF-FBS) [33 pgs.]	
/DB/	3-6	Level 3 Communications, LLC's Objections And Answers To Defendant Limelight Networks, Inc.'s Sixth Set Of Interrogatories, 02/09/2008, in civil action No. 2:07-cv-00589-WDK-FBS [9 pgs.].	
/DB/	3-7	Level 3 Communications, LLC's Objections And Answers To Defendant Limelight Networks, Inc.'s Third Set Of Interrogatories, [redacted], August 18, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [10 pgs].	

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	4-1	Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, [redacted] April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [28 pgs].	
/DB/	4-2	Level 3 Communications, LLC's Second Supplemental Response To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 15), 11/02/2008 in civil action No. 2:07-cv-00589-RGD-FBS [9 pgs.].	
/DB/	4-3	Level 3 Communications, LLC's Second Supplemental Response To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 4) [Redacted], 11/03/2008 in civil action No. 2:07-cv-00589-RGD-FBS [8 pgs.].	
/DB/	4-4	Level 3 Communications, LLC's Supplemental Responses To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 10) [Redacted], June 11, 2008 in civil action No. 2:07-cv-00589-RGD-FBS [6 pgs.].	
/DB/	4-5	Level 3 Communications, LLC's Supplemental Responses To Limelight Networks, Inc.'s First Set Of Interrogatories (Nos. 14 and 15) [Redacted], May 30, 2008 in civil action No. 2:07-cv-00589-RGD-FBS [10 pgs.].	
/DB/	4-6	MEMORANDUM on Markman terms to: Judge Robert G. Doumar, Counsel for Level 3 Communications, Counsel for Limelight from: Court appointed technical expert, Professor Ellen W. Zegura Re: Definition of terms, in civil action No. 2:07-cv-00589-MSD-FBS, Level 3 Communications, LLC v. Limelight Networks, Inc. (August, 2008).	
/DB/	4-7	Order from Judge Doumar in civil action No. 2:07-cv-00589-MSD-FBS, Level 3 Communications, LLC v. Limelight Networks, Inc. instructing Court-Appointed Expert, Professor Ellen W. Zegura Re: Definition of claim terms, July 25, 2008, [2 pgs.].	

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	5-1	Plaintiff Level 3 Communications, LLC's Memorandum Of Law In Opposition To Defendant Limelight Networks, Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, September 30, 2008, Civil Action No. 2:07cv589, [31 pages].	
/DB/	5-2	Response to Summons of May 30, 2008 filed in European Patent Office for European Application No. 99 906 680.6 on October 20, 2008.	
/DB/	5-3	Transcript of 30(b)(6) Deposition of Kevin F. Delgadillo, Tuesday, September 9, 2008 in Civil Action No. 2:07cv589 [39 pages].	
/DB/	5-4	Transcript of Deposition Of Maurice Andrew Collins, September 9, 2008 in Civil Action No. 2:07CV589 [136 pgs.]	
	5-5		
	5-6		
	5-7		

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	1-1	Opinion and Order from U.S. District Judge Mark. S. Davis in Civil Action No. 2:07cv589, denying Defendant LimeLight Networks Inc.'s Motion for Summary Judgment, Norfolk, VA, December 29, 2008.	
	1-2		
	1-3		
	1-4		
	1-5		
	1-6		
	1-7		

Examiner Signature	/Douglas Blair/	Date Considered	01/02/2009

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	1-1	Opinion and Order from U.S. District Judge Mark. S. Davis in Civil Action No. 2:07cv589, denying Defendant LimeLight Networks Inc.'s Motion for Summary Judgment, Norfolk, VA, December 29, 2008.	
	1-2		
	1-3		
	1-4		
	1-5		
	1-6		
	1-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes. If identified, the following is provided. EA = English Abtract, T = Translation, PF = Patent Family.

Electronic Acknowledgement Receipt	
EFS ID:	4538095
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	75948
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	30-DEC-2008
Filing Date:	23-FEB-2005
Time Stamp:	16:22:37
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		PA-IDS_filed_12_30_08.pdf	4346985 3e38872a42643b367073497a558eaece0d1d7d4d	yes	27

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Information Disclosure Statement Letter	1	3
	Information Disclosure Statement (IDS) Filed (SB/08)	4	4
	NPL Documents	5	27
Warnings:			
Information:			
	Total Files Size (in bytes):	4346985	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:		Attorney Docket:	2711-0073
FARBER, David et al.		Group Art Unit:	2142
Application Serial No.:	11/065,412	Examiner:	BLAIR, Douglas B.
Application Filing Date:	February 23, 2005	Confirmation No.:	3296
Title: **SHARED CONTENT DELIVERY INFRASTRUCTURE**		Date:	December 30, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference listed on the attached PTO-1449. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

As Applicant has previously advised the Office, this patent application is related to U.S. Patent Nos. 7,054,935 and 6,654,807, both of which have been asserted in a litigation for patent infringement in a case captioned: *Level 3 Communications, LLC v. Limelight Networks, Inc.* (Civil Action No. 2:07cv589).

A recent document relating to that litigation is cited herewith. Specifically, this IDS cites an Opinion and Order (December 29, 2008) from U.S. District Judge Mark. S. Davis in Civil Action No. 2:07cv589, denying the defendant's Motion for Summary Judgment. The Examiner is again respectfully requested to contact the undersigned with any questions about the related litigation.

This IDS is intended to be in full compliance with the rules, but should the Examiner find any part of its requirement content to have been omitted, prompt notice to that effect is earnestly solicited, along with additional time under Rule 97(f), to enable Applicant to comply fully.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. No certification or fee is required.

Contingent IDS Request Under Rule 97(c): Should a first Office Action on the merits issue with a mailing date which precedes or is the same as the filing date of this IDS, please consider this a Request under Rule 97(c), charge the IDS fee (Rule 17(p)) to our Deposit Account No. 501860 under Order No. 2711-0073 and proceed to consider this IDS under Rule 97(c).

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By: /Brian Siritzky/Reg. No. 37497 Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 ● FAX: (703) 894-6430	



IPW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: December 22, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

As Applicant has previously advised the Office, this patent application is related to U.S. Patent Nos. 7,054,935 and 6,654,807, both of which have been

asserted in a litigation for patent infringement in a case captioned: *Level 3 Communications, LLC v. Limelight Networks, Inc.* (Civil Action No. 2:07-cv-00589). Documents relating to that litigation have been cited in various Information Disclosure Statements, and this IDS includes additional documents from the litigation.

This IDS cites, inter alia, the Court's Claim construction Opinion and Order in Civil Action No. 2:07:cv589, December 10, 2008 [51 pages].

This IDS also cites "Plaintiff Level 3 Communications, LLC's Memorandum Of Law In Opposition To Defendant Limelight Networks, Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, 09/30/08, Civil Action No. 2:07cv589," which includes seven exhibits. The Exhibits are titled and dated as follows:

A	Defendant Limelight Networks, Inc.'s Supplemental Responses to Plaintiff's First set of Interrogatories	04/21/2008
B	Level 3 Communications, LLC's Objections and Responses to Limelight Networks, Inc.'s First Set of Interrogatories.	04/21/2008
C	Defendant Limelight Networks, Inc.'s Second Supplemental Response to Plaintiff's Interrogatory No. 3	08/01/2008
D	Defendant Limelight Networks, Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (Nos. 1-1 1.13-14.17 and 19-20,	09/02/2008
E	Defendant Sockeye Networks, Inc.'s Answer to Complaint and Counterclaims, in Civil Case No. C 02-3708 (EDL)	09/20/2002
F	Deposition Transcript of Livio, Ricciulli, in Civil Action No. 2:07cv589.	09/30/2008
G	List of attorneys of record in Civil Action No. 3:02-cv-03708-CRB	09/20/2002

Civil Action No. 2:07cv589 is scheduled for a two-week trial in Norfolk, Virginia, starting January 5, 2009. The Examiner is again respectfully requested to contact the undersigned with any questions about the related litigation.

This IDS is intended to be in full compliance with the rules, but should the Examiner find any part of its requirement content to have been omitted, prompt notice to that effect is earnestly solicited, along with additional time under Rule 97(f), to enable Applicant to comply fully.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. No certification or fee is required.

Contingent IDS Request Under Rule 97(c): Should a first Office Action on the merits issue with a mailing date which precedes or is the same as the filing date of this IDS, please consider this a Request under Rule 97(c), charge the IDS fee (Rule 17(p)) to our Deposit Account No. 501860 under Order No. 2711-0073 and proceed to consider this IDS under Rule 97(c).

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By: /Brian Siritzky/Reg. No. 37497 Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 • FAX: (703) 894-6430	

OIPE IAP81 DEC 22 2008 PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 5	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	1-1	Amended Answer of Defendant Limelight Networks, Inc., Exhibit 1 to Defendant Limelight Network Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, in Civil Action No. 2:07cv589, September 9, 2008 [18 pgs].	
	1-2	Claim construction Opinion and Order in Civil Action No. 2:07:cv589, December 10, 2008 [51 pages].	
	1-3	Communications from European Patent Office in European Patent Application No. 99 906 680.6: Minutes of Oral Hearing of 11/20/2008, including auxiliary requests; and Invitation under Art. 94(3) and Rule 71(1) EPC (dated Dec. 11, 2008) to amended dependent claims. [32 pgs.]	
	1-4	Declaration Kevin C. Almeroth in Support of Limelight Networks, Inc.'s Motion for Summary Judgment of Non-Infringement and Invalidity of All Patents-In-Suit, filed August 28, 2008 in Civil Action No. 2:07cv589 (RGD-FBS) [39 pgs.].	
	1-5	Defendant Limelight Network Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, in Civil Action No. 2:07cv589, September 9, 2008 [4 pgs].	
	1-6	Defendant Limelight Network Inc.'s Memorandum Of Points And Authorities In Support Of Its Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses (PUBLIC VERSION), in Civil Action No. 2:07cv589, September 9, 2008 [18 pgs.]	
	1-7	Defendant Limelight Network Inc.'s Memorandum of Points and Authorities in Support of its Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, (Public Version, redacted) filed 09/09/208, in civil action No. 2:07-cv-00589-RGD-FBS.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	2-1	Defendant Limelight Network Inc.'s Rebuttal Brief In Support Of Its Combined Motion For Leave To Amend Its Answer And Compel Related Interrogatory Responses (Public version), filed 10/07/2008 in civil action No. 2:07-cv-00589-RGD -FBS.	
	2-2	Defendant Limelight Network Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (Redacted), April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [46 pgs.]	
	2-3	Defendant Limelight Networks, Inc.'s Claim Construction Brief for U.S. Patent Nos. 6,654,807; 7,054,935; and 6,473,405, filed in Civil Action No. 2:07CV589 (RGD-FBS), June 27, 2008 [38 pgs.].	
	2-4	Defendant Limelight Networks, Inc.'s Reply Brief In Support Of Its Motion For Summary Judgment Of Non-Infringement And Invalidity Of All Patents-In-Suit, filed October 1, 2008 in Civil Action No. 2:07cv589 (JBF-FBS) [21 pgs.].	
	2-5	Defendant Limelight Networks, Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (No. 2-4, 12-13, And 19) [Redacted], May 16, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [70 pgs.].	
	2-6	Defendant Limelight Networks, Inc.'s Supplemental Responses To Plaintiff's First Set Of Interrogatories (Nos. 1-11, 13-14, 17 and 19-20, [redacted] 02/09/2008, in civil action No. 2:07-cv-00589-RGD-FBS [101 pgs.].	
	2-7	Exhibit A to Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [68 pgs.].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	3-1	Exhibit B to Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [20 pgs.].	
	3-2	Exhibits A-G of Plaintiff Level 3 Communications, LLC's Memorandum Of Law In Opposition To Defendant Limelight Networks, Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, 09/30/08, Civil Action No. 2:07cv589, (Exhibit A [7 pgs.], B [5 pgs.], C [47 pgs.], D [45 pgs.], E [7 pgs.], F [8 pgs.], G [5 pgs.]).	
	3-3	Exhibits A-V filed in support of Limelight Networks, Inc.'s Motion for Summary Judgment of Non-Infringement and Invalidity of All Patents-In-Suit, August 28, 2008 in Civil Action No. 2:07cv589 (JBF-FBS) [277 pgs.]	
	3-4	Level 3 Communications Supplemental Exhibit A in Civil Action No. 2:07cv589 (JBF -FBS) [69 pgs.]	
	3-5	Level 3 Communications, LLC's Confidential Version Of Its Opposition Brief To Defendant's Motion For Summary Judgment Of Non-Infringement And Invalidity Of All Patents-In-Suit, September 17, 2008, filed in Civil Action No. 2:07cv589 (JBF-FBS) [33 pgs.]	
	3-6	Level 3 Communications, LLC's Objections And Answers To Defendant Limelight Networks, Inc.'s Sixth Set Of Interrogatories, 02/09/2008, in civil action No. 2:07-cv-00589-WDK-FBS [9 pgs.].	
	3-7	Level 3 Communications, LLC's Objections And Answers To Defendant Limelight Networks, Inc.'s Third Set Of Interrogatories, [redacted], August 18, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [10 pgs].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	4-1	Level 3 Communications, LLC's Objections And Responses To Limelight Networks. Inc.'s First Set Of Interrogatories, [redacted] April 21, 2008, in civil action No. 2:07-cv-00589-WDK-FBS [28 pgs].	
	4-2	Level 3 Communications, LLC's Second Supplemental Response To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 15), 11/02/2008 in civil action No. 2:07-cv-00589-RGD-FBS [9 pgs.].	
	4-3	Level 3 Communications, LLC's Second Supplemental Response To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 4) [Redacted], 11/03/2008 in civil action No. 2:07-cv-00589-RGD-FBS [8 pgs.].	
	4-4	Level 3 Communications, LLC's Supplemental Responses To Limelight Networks, Inc.'s First Set Of Interrogatories (No. 10) [Redacted], June 11, 2008 in civil action No. 2:07-cv-00589-RGD-FBS [6 pgs.].	
	4-5	Level 3 Communications, LLC's Supplemental Responses To Limelight Networks, Inc.'s First Set Of Interrogatories (Nos. 14 and 15) [Redacted], May 30, 2008 in civil action No. 2:07-cv-00589-RGD-FBS [10 pgs.].	
	4-6	MEMORANDUM on Markman terms to: Judge Robert G. Doumar, Counsel for Level 3 Communications, Counsel for Limelight from: Court appointed technical expert, Professor Ellen W. Zegura Re: Definition of terms, in civil action No. 2:07-cv-00589-MSD-FBS, Level 3 Communications, LLC v. Limelight Networks, Inc. (August, 2008).	
	4-7	Order from Judge Doumar in civil action No. 2:07-cv-00589-MSD-FBS, Level 3 Communications, LLC v. Limelight Networks, Inc. instructing Court-Appointed Expert, Professor Ellen W. Zegura Re: Definition of claim terms, July 25, 2008, [2 pgs.].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	5-1	Plaintiff Level 3 Communications, LLC's Memorandum Of Law In Opposition To Defendant Limelight Networks, Inc.'s Combined Motion For Leave To Amend Its Answer And To Compel Related Interrogatory Responses, September 30, 2008, Civil Action No. 2:07cv589, [31 pages].	
	5-2	Response to Summons of May 30, 2008 filed in European Patent Office for European Application No. 99 906 680.6 on October 20, 2008.	
	5-3	Transcript of 30(b)(6) Deposition of Kevin F. Delgadillo, Tuesday, September 9, 2008 in Civil Action No. 2:07cv589 [39 pages].	
	5-4	Transcript of Deposition Of Maurice Andrew Collins, September 9, 2008 in Civil Action No. 2:07CV589 [136 pgs.]	
	5-5		
	5-6		
	5-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Document code: WFEE

United States Patent and Trademark Office
Sales Receipt for Accounting Date: 01/30/2009

DWILLIA4 SALE #00000001 Mailroom Dt: 12/22/2008 501860 11065412
01 FC : 1806 180.00 DA

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application Number:	11/065,412	Confirmation Number:	3296
First Named Inventor:	David A. Farber	Examiner:	BLAIR, DOUGLAS B
Filing Date:	02-23-2005	Attorney Docket Number:	2711-0073
Group Art Unit:	2142	Date:	Dec. 10, 2008
Title:	**Shared Content Delivery Infrastructure**		

SUPPLEMENTAL PRELIMINARY AMENDMENT

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Before initial examination, the Applicant respectfully requests the Examiner to include the following as though included in the initial filing.

A listing of the Claims (with claim status) begin on page 2 of this paper.

Remarks begin on page 8 of this paper.

THE CLAIMS:

1-26. (Canceled)

27. (New) A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;

a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level domain name service (DNS) resolution;

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

28. (New) The distributed hosting framework as recited in claim 27, further including a redundant second level name server.

29. (New) A distributed hosting framework operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, the framework comprising:

at least one shared content server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers; and

at least one name server that provides domain name service (DNS) resolution,

wherein resources associated with the first content provider are associated with a first alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the first content provider made from said client machines; and

wherein the resources associated with the second content provider are associated with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the second content provider made from said client machines.

30. (New) The distributed hosting framework as recited in claim 29, wherein the least one shared content server is addressable using the first alias name and the second alias name.

31. (New) The distributed hosting framework as recited in claim 29, further comprising:

a reflector constructed and adapted to associate the resources with the first alias name based, at least in part, on a relative cost of transmitting the resources to the client machines.

32. (New) The distributed hosting framework as recited in claim 29 further comprising:

a set of shared content servers that comprises the at least one content server; and

a reflector constructed and adapted to associate a particular resource with the first alias name after selecting the at least one shared content server from the set of shared content servers in response to a request for the particular resource issued by a particular client machine.

33. (New) The distributed hosting framework as recited in claim 29, wherein at least one of the resources is an embedded object.

34. (New) The distributed hosting framework as recited in claim 29, wherein the requests comprise host tag names that identify specific origin servers from which the resources originate.

35. (New) The distributed hosting framework as recited in claim 34, wherein the host tag names are included in HTTP headers.

36. (New) The distributed hosting framework as recited in claim 32, further comprising:

a load determination mechanism constructed and adapted to characterize load on the set of content servers for use by the reflector in selecting the at least one shared content server.

37. (New) The distributed hosting framework as recited in claim 36, wherein the reflector comprises the load mechanism.

38. (New) The distributed hosting framework as recited in claim 37, wherein the reflector is co-located with an origin server.

39. (New) The distributed hosting framework as recited in claim 36, further comprising:

a cost determination mechanism constructed and adapted to measure cost of delivering resources from each of the set of shared content servers to the particular client machine, the cost determination mechanism for use by the reflector in selecting the at least one shared content server.

40. (New) The distributed hosting framework as recited in claim 39, wherein the reflector comprises the cost determination mechanism.

41. (New) The distributed hosting framework as recited in claim 39, wherein the reflector is co-located with an origin server.

42. (New) The distributed hosting framework as recited in claim 39, wherein the cost comprises information relating a location associated with the particular client machine to each of the set of shared content servers.

43. (New) The distributed hosting framework as recited in claim 29, further including a redundant name server.

44. (New) The distributed hosting framework as recited in claim 29, wherein the at least one name server comprises:

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level DNS resolution.

45. (New) A method, in a distributed hosting framework operative in a computer network for delivering content to client machines and comprising a plurality of origin servers, each of said origin servers having resources associated therewith, at least one name server that provides domain name service (DNS) resolution, and at least one shared content server operable to replicate resources associated with the plurality of origin servers, the method comprising:

associating at least some resources associated with a first origin server of plurality of origin servers with a first alias name that causes the at least one name server to indentify the at least one shared content server to serve the resources associated with the first origin server in response to requests from client machines; and

associating at least some resources associated with a second origin server of said plurality of origin servers with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources associated with the second origin server in response to requests from client machines.

46. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one shared content server from a set of shared

content servers based at least in part on a network cost of delivering resources to the particular client machine.

47. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one content server from a set of shared content servers based at least in part on load characteristics of the set of shared content servers.

48. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one content server from a set of shared content servers based at least in part on a location of the particular client machine relative to the at least one shared content server.

49. (New) The method as recited in claim 45, wherein the resources comprise an HTTP-based resource.

50. (New) The method as recited in claim 45, wherein the resources comprise a File Transfer Protocol (FTP)-based resource.

51. (New) The method of claim 45 wherein the at least one name server comprises:

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level DNS resolution.

REMARKS

By this Amendment, all pending claims have been canceled and new claims 27-52 have been added. Upon entry of this amendment, claims 27-52 will be pending in this application.

This Supplemental Preliminary Amendment is being filed to correctly list the status of all of the claims. In the Preliminary Amendment filed August 5, 2008, Applicant failed to include the status of claims 1-26 as "Canceled". Applicant thanks the Examiner for pointing out this omission and for providing an opportunity to correct it.

LITIGATION AND TRIAL IN RELATED PATENTS

As Applicant has previously advised the Office, this patent application is related to U.S. Patent Nos. 7,054,935 and 6,654,807, both of which have been asserted in a litigation for patent infringement in a case captioned: *Level 3 Communications, LLC v. Limelight Networks, Inc.* (Civil Action No. 2:07-cv-00589). The case was filed on 12/17/2007. Documents relating to that litigation have been cited in various Information Disclosure Statements, and Applicant expects to file additional documents from that litigation.

Civil Action No. 2:07cv589 is scheduled for a two-week trial in Norfolk, Virginia, starting January 5, 2009.

The Examiner is again respectfully requested to contact the undersigned with any questions about the related litigation.

CLARIFICATION ON CLAIM NUMBERING

This application was filed with 24 claims. By a preliminary amendment with the filing of the application, claims 1-24 were canceled and new claim 25 was added. At that time claim 25 was the only pending claim. Then, on June 26,

2007, by another preliminary amendment, 26 new claims were added. These new claims were mistakenly numbered 2-26. Since the previously highest numbered claim was 25, these claims should have been numbered 26-47. However, the Office has entered these 25 claims as claims 2-26. Applicant has therefore numbered the new claims filed herewith as claims 27-52.

For clarification, the new claims filed herewith replace *all* previous claims. Should their numbering be incorrect, the Applicant respectfully requests the Examiner to renumber the claims or to indicate to the Applicant how the claims should be renumbered.

NOTE ON NEW CLAIMS 27 AND 28

New claims 27 and 28 correspond to claims 1 and 3 of **U.S. Patent No. 6,108,703**, Leighton *et al.,* filed May 19, 1999, issued August 22, 2000 ("the '703 Patent"). These claims of the '703 Patent copied herein were previously copied in application no. 09/612,598 on September 18, 2000, prior to one year from the date on which the '703 Patent was granted.

This application is a continuation of 09/612,598 which was a division of 09/021,506, filed February 10, 1998 (now patent no. 6,185,598 – "the '598 Patent"). The specification of this application is the same as that of the '598 Patent (except for the cross reference to related applications in the first paragraph of this application).

Applicant is <u>not</u> requesting an interference with claims 1 and 3 of the '703 Patent since those claims the '703 Patent (corresponding to claims 27 and 28 hereof) have been invalidated in judicial proceedings, as being ***anticipated*** under 35 U.S.C. §102(e) by U.S. Patent No. 6,185,598, from which this application claims priority. *See Akamai Techs. v. Cable & Wireless Internet Services, Inc. et al.,* 344 F.3d 1186, 1197 (Fed. Cir. 2003), attached as Exhibit 1 hereto. As noted by the Federal Circuit in *Akamai Tech.*, "… we hold that claims 1 and 3 of the

'703 patent are anticipated by the '598 patent and are invalid pursuant to 35 U.S.C. § 102." *Id.*

In view of their invalidation, claims 1 and 3 are no longer in the '703 Patent. Furthermore, in view of their invalidation as being *anticipated* by the '598 Patent, it follows that the instant application fully describes and enables those two claims.

CONCLUSION

Applicant respectfully submits that this application is in condition for allowance, and an early action allowing the claims is earnestly solicited. Should the Examiner believe that a telephone call will resolve any outstanding issues in this case, he is invited to telephone the undersigned at the number provided.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/ Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached. **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

Respectfully submitted,

/Brian Siritzky/Reg. No. 37497

By: _____
Brian Siritzky, Ph.D., Reg. No. 37497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 ● FAX: (703) 894-6430

Electronic Acknowledgement Receipt

EFS ID:	4428577
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	75948
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	10-DEC-2008
Filing Date:	23-FEB-2005
Time Stamp:	13:18:40
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Preliminary Amendment	2711_0073_Supplemental_Preliim_Amend.pdf	454901 f8fd7bce95d5d3e1b6f299fd7fe2458a8bc79bd4	no	10

Warnings:	
Information:	

Total Files Size (in bytes):	454901

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (07-06)
Approved for use through 1/31/2007. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/065,412	02/23/2005	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☐		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)	SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
	12/10/2008	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	* 25	Minus	** 48	= 0	X $ =		OR	X $52=	0
	Independent (37 CFR 1.16(h))	* 3	Minus	***5	= 0	X $ =		OR	X $220=	0
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)					
		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =		OR	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =		OR	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".

*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

Legal Instrument Examiner:
/LAMONT MCLAUCHLIN/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: September 23, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. No certification or fee is required.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By: /Brian Siritzky/Reg. No. 37497 Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 ● FAX: (703) 894-6430	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	1-1	Defendant Limelight Networks, Inc.'s Memorandum in Support of its Motion for Summary Judgment of Non-Infringement and Invalidity of all Patents-In-Suit, for Civil Action No. 2:07cv589 (RGD-FBS), August 28, 2008.	
	1-2	Interlocutory Decision in Opposition proceedings for EPO Appln. No. 00 128 346.4 - 2211, Patent No. 1143337, September 22, 2008.	
	1-3	Level 3 Communications, LLC's Non-Confidential Version of its Opposition Brief to Defendant's Motion for Summary Judgment of Non-Infringement and Invalidity of all Patents-In-Suit in Civil Action No. 2:07cv589 (JBF-FBS), Sept. 17, 2008.	
	1-4	Level 3's Memorandum of Law in Reply to Limelight's Claim Construction Brief for U.S. Patent Nos. 6,654,807; 7,054,935; and 6,473,405, in Civil File No. 2:07CV589, June 27, 2008.	
	1-5	Level 3's Memorandum of Law in Reply to Limelight's Claim Construction Brief, in Civil File No. 2:07CV589, July 3, 2008.	
	1-6	Memorandum of Law in Support of Level 3's Proposed Claim Constructions, from Case No. 2:07cv589 RDG FBS, June 13, 2008.	
	1-7	Minutes of Oral Proceedings Appln. No. 00 128 346.4 - 2211, Patent No. 1143337, September 22, 2008.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:		Attorney Docket:	2711-0073
FARBER, David et al.		Group Art Unit:	2142
Application Serial No.:	11/065,412	Examiner:	BLAIR, Douglas B.
Application Filing Date:	February 23, 2005	Confirmation No.:	3296
Title:	**DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date:	September 23, 2008

LETTER / RESPONSE TO OFFICE ACTION

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

On September 19, 2008 the Examiner issued a first Office Action on the merits of the above-identified application. Unfortunately this first Office Action did not consider the currently pending claims (claims 27-51 filed with Applicant's Preliminary Amendment of August 5, 2008). Instead the Examiner considered only a single claim that had been cancelled by earlier amendment.

On September 23, 2008 the undersigned telephoned the Examiner to discuss this matter. The Examiner requested that Applicant file this letter in response to the first Office Action, and that he would then issue a new Office Action on the merits of the claims as filed on August 5, 2008.

Applicant respectfully requests that the Office issue a new action on the merits of the currently pending claims 27-51 filed with Applicant's Preliminary Amendment of August 05, 2008.

Applicant thanks the Examiner for the courtesies extended during the telephone call with the undersigned. The Examiner is kindly requested to contact the undersigned at the number provided with any questions about this matter.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 75948	Respectfully submitted, By: /Brian Siritzky/Reg. No. 37497 Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 ● FAX: (703) 894-6430	

Electronic Acknowledgement Receipt	
EFS ID:	3988542
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	75948
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	23-SEP-2008
Filing Date:	23-FEB-2005
Time Stamp:	14:50:51
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement Letter	PA-IDS-Filed_09_23_09.pdf	166703 088b9feb8e6ca4e2cac7c379a777fc8b1faebaee	no	2

Warnings:

Information:

2	NPL Documents	D103065EPT1Minutes.pdf	514309 bf7d69dbcdaaf752173bb104a0dc68fe0f04e1c9	no	17
Warnings:					
Information:					
3	NPL Documents	D103065EPT1Decision.pdf	1638904 4d64d1bc6b2bc2172442f6536dc5875bfdfe6a61	no	43
Warnings:					
Information:					
4	NPL Documents	Non-ConfidentialBrief.pdf	548910 813892b99f7b69a31478cd50ce76efdca194e6ba	no	33
Warnings:					
Information:					
5	NPL Documents	LLRedactedSJBrief.pdf	1264015 83e997d459e3461b5e8fe3e0066d0d0edd955005	no	31
Warnings:					
Information:					
6	NPL Documents	061308Level3memo_.pdf	8145590 99440dd0483708bd8ba7b8de6f2f33f01aee0328	no	170
Warnings:					
Information:					
7	NPL Documents	Binder2.pdf	1288615 4bb1b7556130b0bd3d2dbd1054acb8f2d1527dd9	no	38
Warnings:					
Information:					
8	Information Disclosure Statement (IDS) Filed (SB/08)	20080923140838581.pdf	174318 abf8a30b78f8466bb992eaacd310dcd7eade0f4d	no	1
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
9	Applicant Arguments/Remarks Made in an Amendment	PA-Letter_re_Office_Action_filed_09_23_08.pdf	312971 0d493149e5170ed35d1afea00b385d561ddf839f	no	2
Warnings:					
Information:					
10	NPL Documents	Binder3.pdf	9863212 9cf8a0bde465456d4eaca073fbc09b82e5ec0c30	no	84

Warnings:	
Information:	
Total Files Size (in bytes):	23917547

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (07-06)
Approved for use through 1/31/2007. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/065,412	02/23/2005	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☐		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT	09/23/2008	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE ($)	ADDITIONAL FEE ($)	OR	OTHER THAN SMALL ENTITY RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 25	Minus	** 48	= 0	X $ =		OR	X $50=	0
	Independent (37 CFR 1.16(h))	* 3	Minus	***5	= 0	X $ =		OR	X $210=	0
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	0

AMENDMENT		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =		OR	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =		OR	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".

*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

Legal Instrument Examiner:
/DIANE WILLIAMS/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
11/065,412	02/23/2005	David A. Farber	2711-0073	3296

75948 7590 09/18/2008

DAVIDSON BERQUIST JACKSON & GOWDEY, LLP
ATTN: BRIAN SIRITZKY, Ph.D.
4300 WILSON BLVD., 7TH FLOOR
ARLINGTON, VA 22203

EXAMINER	
BLAIR, DOUGLAS B	
ART UNIT	**PAPER NUMBER**
2142	

MAIL DATE	DELIVERY MODE
09/18/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	11/065,412	FARBER ET AL.	
	Examiner	Art Unit	
	DOUGLAS B. BLAIR	2142	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *23 February 2005*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *25* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *25* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

> Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

Claim 25 is rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. All of the elements of they system of claim 25 are disclosed as being software elements. Because claim 25 is directed towards software per se it does not fit into any of the statutory categories of invention.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claim 25 is rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Patent Number 5,896,533 to Ramos et al.

As to claim 25, Ramos teaches a system operative in a computer network in which users of client machines connect to a first server, the system comprising: a routine for modifying at least one embedded object URL of a web page to include a host name prepended to a domain

name and path; a set of repeater servers, distinct from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server; and a repeater selector mechanism constructed and adapted to identify, for a particular client machine, an appropriate repeater server from the set of repeater servers; wherein in response to requests for the web page, generated by the client machines, the web page including the modified embedded object URL is served from the first server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the repeater selector mechanism (See Figures 3 and 5 and corresponding text).

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to DOUGLAS B. BLAIR whose telephone number is (571)272-3893. The examiner can normally be reached on 9:00am-5:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Andrew Caldwell can be reached on (571) 272-3868. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Douglas B Blair/
Patent Examiner, Art Unit 2142

Notice of References Cited	Application/Control No. 11/065,412	Applicant(s)/Patent Under Reexamination FARBER ET AL.	
	Examiner DOUGLAS B. BLAIR	Art Unit 2142	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-5,925,106	07-1999	Nielsen, Jakob	709/247
*	B	US-2001/0052024	12-2001	DEVARAKONDA et al.	709/238
*	C	US-6,424,992	07-2002	Devarakonda et al.	709/203
*	D	US-6,061,738	05-2000	Osaku et al.	709/245
*	E	US-5,896,533	04-1999	Ramos et al.	709/217
*	F	US-6,230,196	05-2001	Guenthner et al.	709/223
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	((embed$4 near URL) or plugin) same prepend$3 near3 domain	US-PGPUB; USPAT	OR	OFF	2008/09/15 05:42
L2	77	prepend$3 near3 domain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:43
L3	3	L2 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:45
L4	14	prepend$3 with domain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46
L5	0	L4 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46
L6	0	((embed$4 near URL) or plugin) same prepend$3 near3 domain	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:46
L7	42	((embed$4 near URL) or plugin) same (domain with URL)	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:47
L8	1	L7 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:47
L9	72	subdomain with URL	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:49

L10	5	L9 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:50
L11	0	((embed$4 near URL) or plugin) same subdomain$1	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
L12	18	"embedded object URL"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
L13	1	L12 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 05:55
L14	411	("20040139097" | "20050038851" | "20050114296" | "4591983" | "4594704" | "4726017" | "4803641" | "4839798" | "4847784" | "4920432" | "4922417" | "4943932" | "4949187" | "4949248" | "0449570" | "5029232" | "5130792" | "5132992" | "5136716" | "5172413" | "5191573" | "5253275" | "5253341" | "5287499" | "5287537" | "5291554" | "5341477" | "5371532" | "5410343" | "5414455" | "5442389" | "5442390" | "5442749" | "5471622" | "5475615" | "5508732" | "5515511" | "5519435" | "5528281" | "5539621" | "5542087" | "5544313" | "5544327" | "5550577" | "5550863" | "5550982" | "5557317" | "5572643" | "5590288" | "5592611" | "5594910" | "5603026" | "5619648" | "5623656" | "5625781" | "5627829" | "5630067" | "5633999" | "5634006" | "5638443" | "5644714" | "5646676" | "5649186" | "5659729" | "5666362" | "5671279" | "5675734" | "5682512" | "5699513" | "5712979"	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:00

		"5715453"	"5721914"				
		"5734831"	"5740423"				
		"5742762"	"5751961"				
		"5761507"	"5761663"				
		"5764906"	"5774660"				
		"5774668"	"5777989"				
		"5784058"	"5796952"				
		"5799141"	"5802106"				
		"5802291"	"5812769"				
		"5815664"	"5826031"				
		"5828847"	"5832506"				
		"5832514"	"5835718"				
		"5838906"	"5845303"				
		"5856974"	"5862339"				
		"5867706"	"5867799"				
		"5870546"	"5870559"				
		"5878212"	"5884038"				
		"5890171"	"5893116"				
		"5894554"	"5896533"				
		"5903723"	"5907704"				
		"5913028"	"5913033"				
		"5918010"	"5919247"				
		"5920701"	"5933832"				
		"5935207"	"5945989"				
		"5956489"	"5956716"				
		"5959008"	"5961596"				
		"5966440"	"5968121"				
		"5978791"	"5983214"				
		"5983227"	"5987606"				
		"5991809"	"6002720"				
		"6003030"	"6006264"				
		"6012090"	"6014686"				
		"6014698"	"6018516"				
		"6026440"	"6029175"				
		"6029176"	"6035332"				
		"6038216"	"6038310"				
		"6038610"	"6041324"				
		"6044405"	"6046980"				
		"6049831"	"6052718"				
		"6052730"	"6065051"				
		"6065062"	"6070191"				
		"6081829"	"6092112"				
		"6092204"	"6105028"				
		"6108673"	"6108703"				
		"6112231"	"6112240"				
		"6115357"	"6119143"				
		"6125388"	"6128601"				
		"6128660"	"6134583"				
		"6144375"	"6144702"				
		"6151624"	"6154753"				
		"6163779"	"6167427"				
		"6173311"	"6173322"				
		"6178160"	"6181867"				
		"6185619"	"6189030"				
		"6226642"	"6230196"				
		"6243760"	"6256675"				
		"6266699"	"6269394"				
		"6282569"	"6286045"				

		"6298041" | "6314565" | "6347085" | "6360256" | "6370571" | "6370580" | "6412000" | "6415280" | "6421726" | "6456630" | "6480893" | "6484204" | "6496856" | "6502125" | "6502215" | "6557054" | "6581090" | "6591299" | "6611862" | "6658464" | "6665706" | "6741563" | "6799221" | "6915329" | "6928442" | "7047300").PN.				
L15	212	("20040139097" | "20050038851" | "20050114296" | "4591983" | "4594704" | "4726017" | "4803641" | "4839798" | "4847784" | "4920432" | "4922417" | "4943932" | "4949187" | "4949248" | "0449570" | "5029232" | "5130792" | "5132992" | "5136716" | "5172413" | "5191573" | "5253275" | "5253341" | "5287499" | "5287537" | "5291554" | "5341477" | "5371532" | "5410343" | "5414455" | "5442389" | "5442390" | "5442749" | "5471622" | "5475615" | "5508732" | "5515511" | "5519435" | "5528281" | "5539621" | "5542087" | "5544313" | "5544327" | "5550577" | "5550863" | "5550982" | "5557317" | "5572643" | "5590288" | "5592611" | "5594910" | "5603026" | "5619648" | "5623656" | "5625781" | "5627829" | "5630067" | "5633999" | "5634006" | "5638443" | "5644714" | "5646676" | "5649186" | "5659729" | "5666362" | "5671279" | "5675734" | "5682512" | "5699513" | "5712979" | "5715453" | "5721914" | "5734831" | "5740423" | "5742762" | "5751961" | "5761507" | "5761663" | "5764906" | "5774660" | "5774668" | "5777989" | "5784058" | "5796952" | "5799141" | "5802106" |	US-PGPUB; USPAT	OR	OFF	2008/09/15 06:00

		"5802291"	"5812769"				
		"5815664"	"5826031"				
		"5828847"	"5832506"				
		"5832514"	"5835718"				
		"5838906"	"5845303"				
		"5856974"	"5862339"				
		"5867706"	"5867799"				
		"5870546"	"5870559"				
		"5878212"	"5884038"				
		"5890171"	"5893116"				
		"5894554"	"5896533"				
		"5903723"	"5907704"				
		"5913028"	"5913033"				
		"5918010"	"5919247"				
		"5920701"	"5933832"				
		"5935207"	"5945989"				
		"5956489"	"5956716"				
		"5959008"	"5961596"				
		"5966440"	"5968121"				
		"5978791"	"5983214"				
		"5983227"	"5987606"				
		"5991809"	"6002720"				
		"6003030"	"6006264"				
		"6012090"	"6014686"				
		"6014698"	"6018516"				
		"6026440"	"6029175"				
		"6029176"	"6035332"				
		"6038216"	"6038310"				
		"6038610"	"6041324"				
		"6044405"	"6046980"				
		"6049831"	"6052718"				
		"6052730"	"6065051"				
		"6065062"	"6070191"				
		"6081829"	"6092112"				
		"6092204"	"6105028"				
		"6108673"	"6108703"				
		"6112231"	"6112240"				
		"6115357"	"6119143"				
		"6125388"	"6128601"				
		"6128660"	"6134583"				
		"6144375"	"6144702"				
		"6151624"	"6154753"				
		"6163779"	"6167427"				
		"6173311"	"6173322"				
		"6178160"	"6181867"				
		"6185619"	"6189030"				
		"6226642"	"6230196"				
		"6243760"	"6256675"				
		"6266699"	"6269394"				
		"6282569"	"6286045"				
		"6298041"	"6314565"				
		"6347085"	"6360256"				
		"6370571"	"6370580"				
		"6412000"	"6415280"				
		"6421726"	"6456630"				
		"6480893"	"6484204"				
		"6496856"	"6502125"				
		"6502215"	"6557054"				

		"6581090" | "6591299" | "6611862" | "6658464" | "6665706" | "6741563" | "6799221" | "6915329" | "6928442" | "7047300").PN.				
L16	197	L15 and (@ad<"19980210" @rlad<"19980210")	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:01
L17	12535	709/217,219,245.ccls.	US-PGPUB; USPAT; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/09/15 06:07

9/15/08 6:55:41 AM
C:\Documents and Settings\dblair\My Documents\EAST\workspaces\11065412.wsp



INFORMATION DISCLOSURE STATEMENT BY APPLICANT
FORM PTO-1449 (modified)

Sheet 1 of 7

Application No.	11/065,412
Filing Date	Feb. 23, 2005
First Named Inventor	Farber, David A.
Group Art Unit	2142
Examiner Name	BLAIR, Douglas B.
Attorney Docket No.	2711-0073
Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	1-1	US-5819092	10-06-1998	Ferguson
/DB/	1-2	US-6078943	06-20-2000	Yu
/DB/	1-3	US-6081835	06-27-2000	Antcliff
	1-4			
	1-5			
	1-6			
	1-7			
	1-8			
	1-9			
	1-10			
	1-11			
	1-12			
	1-13			
	1-14			
	1-15			
	1-16			
	1-17			
	1-18			
	1-19			
	1-20			
	1-21			
	1-22			
	1-23			
	1-24			
	1-25			
	1-26			

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	2-1	"Cisco DistributedDirector," Cisco Systems, Inc., White Paper, cited in Expert Report of Kevin Almeroth with date of 1996 (month unknown). [9 pages]	
/DB/	2-2	ALMEROTH, K., et al. "Scalable Delivery of Web Pages Using Cyclic Best-Effort (UDP) Multicast", IEEE INFOCOM, San Francisco, California, USA, June 1998.	
/DB/	2-3	ALMEROTH, K., Exhibit A to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [31 pages].	
/DB/	2-4	ALMEROTH, K., Exhibit B to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [131 pages].	
/DB/	2-5	ALMEROTH, K., Exhibit C to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [157 pages].	
/DB/	2-6	ALMEROTH, K., Exhibit D to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [128 pages].	
/DB/	2-7	ALMEROTH, K., Exhibit E Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [286 pages].	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	3-1	ALMEROTH, K., Exhibit F to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [221 pages].	
/DB/	3-2	ALMEROTH, K., Exhibit G to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [369 pages].	
/DB/	3-3	ALMEROTH, K., Exhibit H to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [59 pages].	
/DB/	3-4	ALMEROTH, K., Exhibit I to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [95 pages].	
/DB/	3-5	ALMEROTH, K., Exhibit J to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [105 pages].	
/DB/	3-6	ALMEROTH, K., Exhibit K-1 to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [152 pages].	
/DB/	3-7	ALMEROTH, K., Exhibit K-2 to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [25 pages].	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	4-1	ALMEROTH, K., Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008. [64 pages and Exhibits A-J, K1, K2]	
/DB/	4-2	BAENTSCH, M., et al., "Introducing Application-Level Replication and Naming into Today's Web," Computer Networks and ISDN Systems, vol. 28, num. 7-11, pgs. 921-930, May 1996.	
/DB/	4-3	Berners-Lee, T. & Connolly, D., "Hypertext Markup Language - 2.0," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1866, Nov. 1995	
/DB/	4-4	BERNERS-LEE, T., et al., "Hypertext Transfer Protocol – HTTP/1.0," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-http-v10-spec-00.ps, March 1995.	
/DB/	4-5	BERNERS-LEE, T., et al., "Hypertext Transfer Protocol – HTTP/1.0," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1945, May 1996.	
/DB/	4-6	Berners-Lee, T.. "Resource Identifiers in WWW: A Unifying Syntax for the Expression of Names and Addresses of Objects on the Network as used in the World-Wide Web," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1630, June 1994.	
/DB/	4-7	Bestavros, A. Demand-based document dissemination to reduce traffic and balance load in distributed information systems. In Proc. IEEE Symp. on Parallel and Distributed Processing, San Antonio, TX, October 1995.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	5-1	CALLON, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1195, December 1990.	
/DB/	5-2	CHANDRA, R. et al., "BGP Communities Attribute," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1997, August 1996.	
/DB/	5-3	CHEN, E., et al., "An Application of the BGP Community Attribute in Multi-home Routing," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1998, August 1996.	
/DB/	5-4	COLAJANNI, M., et al., "Scheduling Algorithms for Distributed Web Servers," International Conf. on Distributed Computing Systems (ICDCS), Baltimore, Maryland, USA, May 1997.	
/DB/	5-5	Communication from European Patent Office dated May 30, 2008 in patent application EP 99906680 (12 pages).	
/DB/	5-6	ERIKSSON, H., "MBONE: The Multicast Backbone," Communications of the ACM, vol. 37, num. 8, Pg. 54-60, August 1994.	
/DB/	5-7	FIELDING, R., et al., "Hypertext Transfer Protocol – HTTP/1.1," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-http-v11-spec-00.txt, November 1995.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 6 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	6-1	FIELDING, R., et al., "Hypertext Transfer Protocol – HTTP/1.1," Internet Engineering Task Force(IETF) Request for Comments (RFC) 2068, January 1997.	
/DB/	6-2	MALKIN, G., "RIP Version 2 Carrying Additional Information," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1388, January 1993.	
/DB/	6-3	MOY, J., "OSPF Version 2," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1583, March 1994.	
/DB/	6-4	PARRIS C., et al, "A Dynamic Connection Management Scheme for Guaranteed Performance Services in Packet-Switching Integrated Services Networks," UC Berkeley Computer Science Division Tech. Report TR-93-005, 1993.	
/DB/	6-5	PARRIS C., et al, "The Dynamic Management of Guaranteed Performance Connections in Packet Switched Integrated Service Networks," UC Berkeley Computer Science Division and International Computer Science Institute Tech. Report CSD-94-859, 1994.	
/DB/	6-6	PARTRIDGE, C., et al., "Host Anycasting Service," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1546, November 1993.	
/DB/	6-7	POVEY, D. et al., "A distributed internet cache," in Proc. of the 20th Australasian Computer Science Conf., Feb. 1997.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 7 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES			
Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	7-1	REKHTER Y., et al., "A Border Gateway Protocol 4 (BGP-4)," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1771, March 1995.	
/DB/	7-2	REKHTER Y., et al., "Application of the Border Gateway Protocol in the Internet," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1772, March 1995.	
/DB/	7-3	TOPOLCIC, C., "Experimental Internet Stream Protocol, Version 2 (ST-II)," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1190, October 1990.	
/DB/	7-4	TOUCH, J. et al., "The X-Bone,". Third Global Internet Mini-Conference at Globecom '98. Sydney, Australia, Nov. 1998, pp. 59-68 (pp. 44-52 of the mini-conference).	
/DB/	7-5	TRAINA, P., "BGP-4 Protocol Analysis," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1774, March 1995.	
/DB/	7-6	TRAINA, P., "Experience with the BGP-4 protocol," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1773, March 1995.	
/DB/	7-7	WESSELS, D. et al., RFC 2187, "Application of Internet Cache Protocol (ICP), Version 2", Sep. 1997, Memo; Internet RFC/STD/FYI/BCP Archives; Network Working Group.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.



INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 1	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
/DB/	1-1	Complaint for Injunctive and Other Relief, filed 12/17/2007 in Civil Action No. 2:07CV589	
/DB/	1-2	Report on the Filing or Determination of an Action Regarding a Patent or Trademark, dated 12/28/2007, in connection with Civil Action No. 2:07CV589 filed 12/17/2007.	
/DB/	1-3	Rule 7.1 Financial Interest Disclosure Statement filed 12/17/2007 in Civil Action No. 2:07CV589	
	1-4		
	1-5		
	1-6		
	1-7		
	1-8		
	1-9		
	1-10		

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:	
DAVID A. FARBER, ET AL.	Art Group:
Application No.: Unassigned	Examiner:
Filed: February 23, 2005	
For: **Optimized Network Resource Location**	

INFORMATION DISCLOSURE STATEMENT UNDER 37 C.F.R. §1.97

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

In accordance with the duty of disclosure, enclosed is a copy of IDS Citation Form PTO/SB/08 or PTO-1449, together with copies of the documents cited on that form, except for copies not required to be submitted (e.g., copies of U.S. patents and U.S. published patent applications need not be enclosed). This IDS and IDS Citation Form are being submitted concurrently with the Continuation Application. It is respectfully requested that the cited references be considered and that the enclosed copy of PTO/SB/08 be initialed by the Examiner to indicate such consideration and a copy thereof returned to applicant(s). Some or all of the references listed on the enclosed PTO/SB/08 were previously identified in the parent application (Application No. 09/612,598, filed July 7, 2000) and copies of the references were furnished at that time. Accordingly, per 37 CFR §1.98(d)(1) additional copies of those references are not submitted herewith.

The submission of this Information Disclosure Statement is not to be construed as a representation that a search has been made in the subject application and is not to be construed as an admission that the information cited in this statement is material to patentability.

Please charge any fees due to Deposit Account 02-2666. A duplicate copy of the Fee Transmittal (PTO/SB/17) is enclosed for this purpose.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: February 23, 2005

William W. Schaal, Reg. No. 39,018

12400 Wilshire Boulevard, 7th Floor
Los Angeles, CA 90025
Telephone: (714) 557-3800

I hereby certify that this correspondence is being deposited with the United States Postal Service on the date shown below with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Susan McFarlane — *02-23-05*

Susan McFarlane — *Date*

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	
Filing Date	February 23, 2005
First Named Inventor	Farber
Group Art Unit	
Examiner Name	

Sheet 1 of 7

Attorney Docket No: 6233P001DC3

US PATENT DOCUMENTS

Examiner Initial *	USP Document Number	Publication Date	Name of Patentee or Applicant of cited Document	Class	Subclass	Filing Date If Appropriate
/DB/	US-4,495,570	01/22/1985	Kitajima, Hiroyuki , et al.			01/07/1982
/DB/	US-4,594,704	06/10/1986	Ollivier, Yves			11/16/1984
/DB/	US-4,726,017	02/16/1988	Krum, Harry A., et al.			05/21/1985
/DB/	US-4,839,798	06/13/1989	Eguchi, Hiroyoshi, etal.			11/06/1985
/DB/	US-4,920,432	04/24/1990	Eggers, Derek C., et al.			01/12/1998
/DB/	US-4,922,417	05/01/1990	Churm, Brian R., et al.			10/24/1986
/DB/	US-4,949,187	08/14/1990	Cohen, Jason M.			12/16/1988
/DB/	US-4,949,248	08/14/1990	Caro, Marshall A.			07/15/1988
/DB/	US-5,/26,031	10/20/1998	Nielsen, Jakob			06/10/1996
/DB/	US-5,172,413	12/15/1992	Bradley, Graham C., et al.			12/20/1990
/DB/	US-5,253,341	10/12/1993	Rozmanith, Anthony I., et al.			04/11/1991
/DB/	US-5,287,499	02/15/1994	Nemes, Richard M.			05/16/1991
/DB/	US-5,287,537	02/15/1994	Newmark, Rona J., et al.			09/21/1992
/DB/	US-5,291,554	03/01/1994	Morales,			08/19/1992
/DB/	US-5,341,477	08/23/1994	Pitken, Richard P., et al.			08/06/1993
/DB/	US-5,371,532	12/06/1994	Gleman, Alexander , et al.			05/15/1992
/DB/	US-5,410,343	04/25/1995	Coddington, Carl D., et al.			04/02/1993
/DB/	US-5,414,455	05/09/1995	Hooper, Donald F., et al.			07/07/1993
/DB/	US-5,442,389	08/15/1995	Blahut, Donald E., et al.			12/28/1992
/DB/	US-5,442,390	08/15/1995	Hooper, Donald F., et al.			07/07/1993
/DB/	US-5,471,622	11/28/1995	Eadline, Douglas J.			04/20/1995
/DB/	US-5,475,615	12/12/1995	Lin, K D.			12/23/1993
/DB/	US-5,508,732	04/16/1996	Bottomley, John F., et al.			03/22/1993
/DB/	US-5,515,511	05/07/1996	Nguyen, Tam M., et al.			06/06/1994
/DB/	US-5,519,435	05/21/1996	Anderson, Michael H.			09/01/1994
/DB/	US-5,528,281	06/18/1996	Grady, John , et al.			07/30/1993
/DB/	US-5,539,621	07/23/1996	Kikinis, Dan			06/23/1995
/DB/	US-5,542,087	07/30/1996	Neimat, Marie-Anne , et al.			10/15/1993
/DB/	US-5,544,313	08/06/1996	Shachnai, Hadas , et al.			05/11/1994
/DB/	US-5,544,327	08/06/1996	Dan, Asit , et al.			07/26/1994
/DB/	US-5,550,577	08/27/1996	Verbiest, Willem J., et al.			05/19/1994
/DB/	US-5,550,863	08/27/1996	Yurt, Paul , et al.			10/08/1993
/DB/	US-5,550,982	08/27/1996	Long, James E., et al.			06/24/1993
/DB/	US-5,557,317	09/17/1996	Nishio, Makoto , et al.			05/22/1995
/DB/	US-5,572,643	11/05/1996	Judson, David H.			10/19/1995
/DB/	US-5,590,288	12/31/1996	Castor, Patrick F., et al.			03/13/1995

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	
Filing Date	February 23, 2005
First Named Inventor	Farber
Group Art Unit	
Examiner Name	

Sheet 2 of 7

Attorney Docket No: 6233P001DC3

/DB/	US-5,592,611	01/04/1997	Midgely, Christopher W., et al.			03/14/1995
/DB/	US-5,594,910	01/14/1997	Filepp, Robert , et al.			11/26/1993
/DB/	US-5,603,026	02/11/1997	Demers, Alan J., et al.			12/07/1994
/DB/	US-5,619,648	04/08/1997	Canale,			
/DB/	US-5,623,656	04/22/1997	Lyons, Kenneth B.			12/15/1994
/DB/	US-5,625,781	04/29/1997	Cline, Troy L., et al.			10/31/1995
/DB/	US-5,627,829	05/06/1997	Gleeson, Bryan J., et al.			06/06/1995
/DB/	US-5,633,999	05/27/1997	Clowes, Richard F., et al.			12/18/1995
/DB/	US-5,634,006	05/27/1997	Baugher, Mark J., et al.			10/11/1995
/DB/	US-5,638,443	06/10/1997	Stefik, Mark J., et al.			11/23/1994
/DB/	US-5,644,714	07/01/1997	Kikinis, Dan			04/10/1995
/DB/	US-5,646,676	05/30/1995	Dewkett, Thomas J., et al.			05/30/1995
/DB/	US-5,666,362	09/09/1997	Chen, Cheng T., et al.			07/25/1995
/DB/	US-5,671,279	09/23/1997	Elgamal, Taher			11/13/1995
/DB/	US-5,682,512	10/28/1997	Tetrick, R. S.			06/30/1995
/DB/	US-5,699,513	12/16/1997	Feigen, Ronald G., et al.			03/31/1995
/DB/	US-5,712,979	01/27/1998	Graber, Terry E., et al.			09/20/1995
/DB/	US-5,715,453	02/03/1998	Stuart, Gordon G., et al.			05/31/1996
/DB/	US-5,721,914	02/24/1998	DeVries, Larry R.			09/14/1995
/DB/	US-5,734,831	03/31/1998	Sanders, James B.			04/26/1996
/DB/	US-5,740,423	04/14/1998	Logan, James R., et al.			04/14/1998
/DB/	US-5,742,762	04/21/1998	Scholl, Thomas H., et al.			05/19/1995
/DB/	US-5,751,961	05/12/1998	Smyk, Darek A.			01/31/1996
/DB/	US-5,761,507	06/02/1998	Govett, Ian R.			03/05/1996
/DB/	US-5,761,663	06/02/1998	Lagarde, Konrad C., et al.			05/15/1997
/DB/	US-5,774,660	06/30/1998	Brendel, Juergen , et al.			08/05/1996
/DB/	US-5,774,668	06/30/1998	Choquier, Philippe , et al.			06/07/1995
/DB/	US-5,777,989	07/07/1998	McGarvey, John R.			12/19/1995
/DB/	US-5,784,058	07/21/1998	LaStrange, Thomas E., et al.			05/28/1996
/DB/	US-5,796,952	08/18/1998	Davis, Owen , et al.			03/21/1997
/DB/	US-5,799,141	08/25/1998	Galipeau, Kenneth J., et al.			06/09/1995
/DB/	US-5,802,291	09/01/1998	Balick, Maurice , et al.			03/30/1995
/DB/	US-5,812,769	09/22/1998	Graber, Terry E., et al.			09/20/1995
/DB/	US-5,828,847	10/27/1998	Gehr, Chuck R., et al.			04/19/1996
/DB/	US-5,832,506	11/03/1998	Kuzma, Andrew J.			03/29/1996
/DB/	US-5,835,718	11/10/1998	Blewett, Charles D.			04/10/1996
/DB/	US-5,838,906	11/17/1998	Doyle, Michael D., et al.			10/17/1994
/DB/	US-5,856,974	01/05/1999	Gervais, Joseph L., et al.			01/05/1999
/DB/	US-5,862,339	01/19/1999	Bonnaure, Patrick P.			07/09/1996

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	
Filing Date	February 23, 2005
First Named Inventor	Farber
Group Art Unit	
Examiner Name	

Sheet 3 of 7

Attorney Docket No: 6233P001DC3

/DB/	US-5,867,706	02/02/1999	Martin, Sean J., et al.			02/02/1999
/DB/	US-5,867,799	02/02/1999	Lang, Andrew K., et al.			04/04/1996
/DB/	US-5,870,546	02/09/1999	Kirsch, Steven T.			12/23/1997
/DB/	US-5,870,559	02/09/1999	Leshem, Eran , et al.			04/11/1997
/DB/	US-5,878,212	03/02/1999	Civanlar, Seyhan , et al.			07/31/1995
/DB/	US-5,884,038	03/16/1999	Kapoor, Rupesh			05/02/1997
/DB/	US-5,890,171	03/30/1999	Blumer, Thomas P., et al.			08/06/1996
/DB/	US-5,896,533	04/20/1999	Ramos, Daniel , et al.			07/06/1995
/DB/	US-5,903,723	05/11/1999	Beck, Robert D., et al.			12/21/1995
/DB/	US-5,913,028	06/15/1999	Wang, Frank , et al.			10/06/1995
/DB/	US-5,913,033	06/15/1999	Grout, Gordon B.			12/20/1996
/DB/	US-5,918,010	06/29/1999	Appleman, Kenneth H., et al.			02/06/1998
/DB/	US-5,919,247	07/06/1999	Van Hoff, Arthur , et al.			07/24/1996
/DB/	US-5,933,832	08/03/1999	Suzuoka, Takashi , et al.			09/10/1997
/DB/	US-5,945,989	08/31/1999	Freishtat, Gregg S., et al.			03/25/1997
/DB/	US-5,956,716	09/21/1999	Kenner, Brian , et al.			06/07/1996
/DB/	US-5,958,008	09/28/1999	Pogrebisky, Michael , et al.			04/11/1997
/DB/	US-5,961,596	10/05/1999	Takubo, Shunji , et al.			02/14/1997
/DB/	US-5,968,121	10/19/1999	Logan, Keith W., et al.			08/13/1997
/DB/	US-5,983,214	11/09/1999	Lang, Andrew K., et al.			11/05/1998
/DB/	US-5,983,227	11/09/1999	Nazem, Farzad , et al.			06/12/1997
/DB/	US-5,991,809	11/23/1999	Kriegsman, Mark E.			07/25/1997
/DB/	US-6,006,264	12/21/1999	Colby, Steven , et al.			03/30/1998
/DB/	US-6,012,090	01/04/2000	Chung, Pi-yu , et al.			03/14/1997
/DB/	US-6,029,175	02/22/2000	Chow, Yen-whei , et al.			06/07/1996
/DB/	US-6,035,332	03/07/2000	Ingrassia, Jr., Michael I., et al.			10/06/1997
/DB/	US-6,038,310	03/14/2000	Hollywood, Steven , et al.			08/01/1995
/DB/	US-6,041,324	03/21/2000	Earl, Joel R., et al.			11/17/1997
/DB/	US-6,052,718	04/18/2000	Gifford, David K.			01/07/1997
/DB/	US-6,052,730	04/18/2000	Feliciano, Ramon M., et al.			01/09/1998
/DB/	US-6,065,051	05/16/2000	Steele, Douglas W., et al.			04/15/1998
/DB/	US-6,065,062	05/16/2000	Periasamy, Ravi , et al.			12/10/1997
/DB/	US-6,070,191	05/30/2000	Narendran, Balakrishnan , et al.			10/17/1997
/DB/	US-6,081,829	06/27/2000	Sidana, Ashmeet S.			01/31/1996
/DB/	US-6,092,112	07/18/2000	Fukushige, Yoshio			01/14/1997
/DB/	US-6,108,673	08/22/2000	Brandt, Marcia L., et al.			02/25/1997
/DB/	US-6,108,703	08/22/2000	Leighton, F. T., et al.			05/19/1999

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	
	Filing Date	February 23, 2005
	First Named Inventor	Farber
	Group Art Unit	
	Examiner Name	
Sheet 4 of 7	Attorney Docket No: 6233P001DC3	

	US-6,112,231	08/29/2000	DeSimone, Antonio , et al.			10/27/1998
/DB/	US-6,112,240	08/29/2000	Pogue, Michael A., et al.			09/03/1997
/DB/	US-6,128,601	10/03/2000	Van Horne, Peter , et al.			03/02/1998
/DB/	US-6,128,660	10/03/2000	Grimm, Steven M., et al.			03/20/1997
/DB/	US-6,144,375	11/07/2000	Jain, Ramesh C., et al.			08/14/1998
/DB/	US-6,163,779	12/19/2000	Mantha, Murthy L., et al.			09/29/1997
/DB/	US-6,178,160	01/23/2001	Bolton, Derek W., et al.			12/23/1997
/DB/	US-6,185,619	02/06/2001	Joffe, Rodney L., et al.	B1		11/07/1997
/DB/	US-6,189,030	02/13/2001	Kirsch, Steven T., et al.			05/01/1998
/DB/	US-6,226,642	05/01/2001	Beranek, Michael J., et al.			09/11/1997

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Foreign Document No	Publication Date	Name of Patentee or Applicant of cited Document	Class	Subclass	T[2]
/DB/	CA-2202572	10/14/1998	Law, Ka LUn E., et al.			
/DB/	EP-0801487A2	10/15/1997	Charles, Douglas B.			
/DB/	EP-0824236A2	02/18/1998	Baird, Richard , et al.			
/DB/	EP-0865180A2	09/16/1998	Chung, Pi-Yu , et al.			
/DB/	GB-2281793	03/15/1995	Wallis, Graham D., et al.			
/DB/	WO-97/11429	03/27/1997	Graber, Terry E., et al.			
/DB/	WO-97/29423	08/14/1997	Stanford-Clark, Andrew , et al.			
/DB/	WO-98/04985	02/05/1998	Kriegsman, Mark			
/DB/	WO-98/06033	02/12/1998	Agranat, Ian D., et al.			

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No [1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/DB/		"Cache Director System", http:/www.cisco.com/warp/public/751/cache/cds ov.htm, Cisco Systems, Inc.,(1997),1-5	
/DB/		"Cache Director Systems Technology Overview", http:/www.cisco.com/warp/public/751/cache/cds ds.htm, Cisco System, Inc.,(1997),1-3	
/DB/		"Cisco Cache Engine", http:/www.cisco.com/warp/public/751/cache/index.html, Cisco Systems, Inc.,	
/DB/		"Cisco Introduces Web Cache Product for Scaling the Internet", http:/www.cisco.com/warp/public;146/1947.html, Cisco Systems, Inc.,(1997),	
/DB/		"Cisco Systems Ships Cisco LocalDirector to Meet Demands of High-volume Web Traffic", http://www.cisco.com/wrap/public/146/323.html, Cisco Systems, Inc.,(1996),1-2	
/DB/		"Cisco Take Global Route", PC Week News, (2/17/1997),23	
/DB/		"Communication (1 page) and European Search Report for", Application No.	

EXAMINER	/Douglas Blair/	DATE CONSIDERED	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	
Filing Date	February 23, 2005
First Named Inventor	Farber
Group Art Unit	
Examiner Name	

Sheet 5 of 7

Attorney Docket No: 6233P001DC3

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No [1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/DB/		EP 00 12 8346 (3 pages), (08/24/2001),	
/DB/		"Excite Chooses Cisco Local Director to Support its Growing Service", http:/www.cisco.com/wrap/public/146/1857.html, Cisco Systems, Inc.,(1997),1-2	
/DB/		"Global IP/PX Service Should Keep Network Delays Down", Infoworld, (01/20/1997),1 page	
/DB/		"How to Cost-Effectively Scale Web Servers", http://www.cisco.com/warp/public/784//5.html, Cisco Systems, Inc.,(1996),1-5	
/DB/		"http://www.cisco.com/warp/public/751/distdir/dd wp.htm", Cisco Systems Inc., (1997),1-15	
/DB/		"http://www.cisco.com/wrap/public/751/lodir/lodir ds.htm", Cisco System Inc., (1997),1-4	
/DB/		"International Search Report for PCT/US99/01477", (06/08/1999),	
/DB/		"Local Area Network Server Replacement Procedure", IBM Technical Disclosure Bulletin, Vol. 38, No. 1,(January 1995),235-236	
/DB/		"Local Director", http://.cisco.com/warp/public/751/lodir/index.html, 2 pages	
/DB/		"Local Director", Cisco Systems, Inc., 1-3	
/DB/		"Patent Abstracts of Japan, Electronic Mail Multiplexing System and Communication Control Method in the System", JP 05162529, (06/30/1993),	
/DB/		"Patent Abstracts of Japan, Method and Device for Repeating and Converting Information", JP 08328583, (09/1996),	
/DB/		"Scaling the Internet News Service", http:/www.cisco.com/warp/public/751/lodir/news wp.htm, Cisco Systems, Inc.,(1997),1-11	
/DB/		"Scaling the World Wide Web", http://www.cisco.com/warp/public/751/lodir/swww wp.html, Cisco Systems, Inc.,(1996),1-7	
/DB/		ADLER, "Distributed Coordination Models for Client/Server Computing", Computer, Vol. 28, No. 4,(04/01/1995),14-22	
/DB/		ANDRESEN, et al., "SWEB: Towards a Scalable World Wide Web Server on Multicomputers", Proceedings of IPPS, (04/15/1996),850-856	
/DB/		BESTAVROS, AZER , "Speculative Data Dissermination and Service to Reduce Server Load, Network Traffic and Service Time in Distributed Information Systems", In Proceedings of ICDE '96: The 1996 International Conference on Data Engineering, (March 1996),4 pages	
/DB/		CARTER, J. LAWRENCE , et al., "Universal Classes of Hash Functions", Journal of Computer and System Sciences, Vol. 18, No. 2, (April 1979),143-154	
/DB/		CHANKHUNTHOD, ANAWAT , et al., "A Hierarchical Internet Object Cache", USENIX Proceedings, (January 1996),153-163	
/DB/		CORMEN, THOMAS H., et al., "Introduction to Algorithms", The MIT Press, Cambridge, Massachusetts, (1994),219-243, 991-993	
/DB/		DE BRA, P.M.E. , et al., "Information Retrieval in the World Wide Web: Making	

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional) [2] Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	
	Filing Date	February 23, 2005
	First Named Inventor	Farber
	Group Art Unit	
	Examiner Name	
Sheet 6 of 7	Attorney Docket No: 6233P001DC3	

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS			
Examiner Initials*	**Cite No [1]**	**Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.**	**T[2]**
/DB/		client-based searching feasible", Computer Networks and ISDN Systems, NL, North Holland Publishing, Amsterdam, Vol. 27. No. 2, ISSN: 0169-7552, (11/01/1994),183-192	
/DB/		DEERING, STEPHEN , et al., "Multicast Routing in Datagram Internetworks and Extended LANs", ACM Transactions on Computer Systems, Vol. 8, No. 2, (May 1990),85-110	
/DB/		DEVINE, ROBERT , "Design and Implementation of DDH: A Distributed Dynamic Hashing Algorithm", In Proceedings of 4th International Conference on Foundations of Data Organizations and Algorithms, (1993),101-114	
/DB/		FEELEY, MICHAEL , et al., "Implementing Global Memory Management in a Workstation Cluster", In Proceedings of the 15th ACM Symposium on Operating Systems Principles, (1995),201-212	
/DB/		FLOYD, SALLEY , et al., "A Reliable Multicast Framework for Light-Weight Sessions and Application Level Framing", In Proceeding of ACM SIGCOMM '95, 342-356	
/DB/		FREDMAN, MICHAEL , et al., "Storing a Sparse Table with 0(1) Worst Case Access Time", Journal of the Association for Computing Machinery, Vol. 31, No. 3, (July 1984),538-544	
/DB/		GRIGNI, MICHELANGELO , et al., "Tight Bounds on Minimum Broadcasts Networks", SIAM Journal of Discrete Mathematics, Vol. 4, No. 2, (May 1991),207-222	
/DB/		GWERTZMAN, JAMES , et al., "The Case for Geographical Push-Caching", Technical Report HU TR 34-94 (excerpt), Harvard University, DAS, Cambridge, MA 02138, (1994),2 pages	
/DB/		GWERTZMAN, J.S. , et al., "The case for Geographical push-caching", Workshop on Hot Topics in Operating Systems, (05/04/1995),51-55	
/DB/		GWERTZMAN, JAMES , et al., "World-Wide Web Cache Consistency", In Proceedings of the 1996 USENIX Technical Conference, (January 1996),8 pages	
/DB/		KARGER, DAVID , et al., "Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web", In Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, (May 1997),654-663	
/DB/		LITWIN, WLTHOLD , et al., "LH-A Scaleable, Distributed Data Structure", ACM Transactions on Database Systems, Vol. 21, No. 4, (December 1996),480-525	
/DB/		MALPANI, RADHIKA , et al., "Making World Wide Web Caching Servers Cooperate", In Proceedings of World Wide Web Conference, (1996),6 pages	
/DB/		MOURAD, et al., "Scalable Web Server Architectures", Proceedings IEEE Symposium on Computers and Communications, (07/01/1997),12-16	
/DB/		NAOR, MONI , et al., "The Load, Capacity and Availability of Quorum Systems", In Proceedings of the 35th IEEE Symposium on Foundations of Computer	

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional) [2] Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	
	Filing Date	February 23, 2005
	First Named Inventor	Farber
	Group Art Unit	
	Examiner Name	
Sheet 7 of 7	Attorney Docket No: 6233P001DC3	

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS			
Examiner Initials*	**Cite No [1]**	**Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.**	**T[2]**
/DB/		Science, (November 1994),214-225	
/DB/		NISAN, NOAM , "Pseudorandom Generations for Space-Bounded Computation", In Proceedings of the Twenty-Second Annual ACM Symposium on Therory of Computing, (May 1990),204-212	
/DB/		PALMER, MARK , et al., "Fido: A Cache that Learns to Fetch", In Proceedings of the 17th International Conference on Very Large Data Bases, (September 1991),255-264	
/DB/		PANIGRAHY, RINA , "Relieving Hot Spots on the World Wide Web", Massachusetts Institute of Technology, (June 1997),1-66	
/DB/		PELEG, DAVID , et al., "The Availability of Quorum Systems", Information and Computation, 123, (1995),210-223	
/DB/		PLAXTON, C. GREG , et al., "Fast Fault-Tolerant Concurrent Access to Shared Objects", In Proceedings of the 37th IEEE Symposium of Foundations of Computer Science, (1996),570-579	
/DB/		RABIN, MICHAEL , "Efficient Dispersal of Information for Security, Load Balancing and Fault Tolerance", Journal of the ACM, Vol. 36, No. 2, (April 1989),335-348	
/DB/		RAVI, R. , "Rapid Rumor Ramification: Approximating the Minimum Broadcast Time", In Proceedings of the 35th IEEE Symposium on Foundation of Computer Science, (November 1994),202-213	
/DB/		SCHMIDT, JEANETTE , et al., "Chernoff-Hoeffding Bounds for Applications with Limited Independence", In Proceedings of the 4th ACS-SIAM Symposium on Discrete Algorithms, (1993),331-340	
/DB/		TARJAN, ROBERT E., et al., "Storing a Sparse Table", Communications of the ACM, Vol. 22, No. 11, (November 1979),606-611	
/DB/		VITTER, JEFFREY S., et al., "Optimal Prefetching via Data Compression", In Proceedings of 32nd IEEE Symposium on Foundations of Computer Science, (November 1991),121-130	
/DB/		WEGMAN, MARK , et al., "New Hash Functions and Their Use in Authentication and Set Equality", Journal of Computer and System Sciences, Vol. 22, (June 1981),265-279	
/DB/		YAO, ANDREW C., "Should Tables be Sorted?", Journal of Association for Computing Machinery, Vol. 28, No. 3, (July 1981),615-628	

EXAMINER	/Douglas Blair/	**DATE CONSIDERED**	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

OIPE
MAY 12 2008
PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
FORM PTO-1449 (modified)

Sheet 1 of 5

Application No.	11/065,412
Filing Date	Feb. 23, 2005
First Named Inventor	Farber, David A.
Group Art Unit	2142
Examiner Name	BLAIR, Douglas B.
Attorney Docket No.	2711-0073
Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	1-1	US-5708780	01/13/98	Levergood et al.
/DB/	1-2	US-5826031	10/1998	Nielsen
/DB/	1-3	US-5838906	11/1998	Doyle et al.
/DB/	1-4	US-5931904	1999-08-03	Banga
/DB/	1-5	US-5944780	1999-08-31	Chase
/DB/	1-6	US-5973696	10-26-99	Agranat, et al.
/DB/	1-7	US-5996025	11/30/99	Day
/DB/	1-8	US-6006264	12/1999	Colby et al.
/DB/	1-9	US-6021426	2000-02-01	Douglis
/DB/	1-10	US-6041307	March 21, 2000	Ahuja, et al.
/DB/	1-11	US-6098078	2000-08-01	Gehani
/DB/	1-12	US-6108703	08/2000	Leighton et al.
/DB/	1-13	US-6115752	09/2000	Chauhan
/DB/	1-14	US-6125394	2000-09-26	Rabinovich
/DB/	1-15	US-6144996	11/2000	Starnes et al.
/DB/	1-16	US-6154738	11/2000	Call
/DB/	1-17	US-6154777	11/28/2000	Ebrahim
/DB/	1-18	US-6185598	02/2001	Farber
/DB/	1-19	US-6189039	02/13/01	Harvey
/DB/	1-20	US-6195680	02/27/01	Goldszmidt
/DB/	1-21	US-6226618	05/01/01	Downs
/DB/	1-22	US-6263313	07/17/01	Milsted
/DB/	1-23	US-6272566	08/07/01	Craft
/DB/	1-24	US-6282574	08/28/2001	Voit
/DB/	1-25	US-6311214	10/2001	Rhoads
/DB/	1-26	US-6330602	Dec-11-01	Law et al.

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS				
Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	2-1	US-6332195	12/2001	Green et al.
/DB/	2-2	US-6398245	06/04/02	Gruse
/DB/	2-3	US-6418421	07/09/02	Hurtado
/DB/	2-4	US-6442549	08/2002	Schneider
/DB/	2-5	US-6460082	10/01/02	Lumelsky
/DB/	2-6	US-6463454	10/08/02	Lumelsky
/DB/	2-7	US-6463508	10/08/02	Wolf
/DB/	2-8	US-6470389	Oct-22-02	Chung et al
/DB/	2-9	US-6587837	07/01/03	Spagna
/DB/	2-10	US-6654807	25-Nov-03	Farber et al.
/DB/	2-11	US-6763377	07/13/04	Belknap et al
/DB/	2-12	US-6859791	02/22/05	Spagna
/DB/	2-13	US-6963910	11/08/05	Belknap
/DB/	2-14	US-6973485	12/2005	Ebata et al.
/DB/	2-15	US-7103564	09/05/06	Ehnebuske
/DB/	2-16	US-7110984	09/19/06	Spagna
/DB/	2-17	US-7117259	10/03/06	Rohwer
/DB/	2-18	US-7188085	03/06/07	Pelletier
/DB/	2-19	US-7206748	04/17/07	Gruse
	2-20			
	2-21			
	2-22			
	2-23			
	2-24			
	2-25			
	2-26			

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 5	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication Date	Name of Patentee or Applicant of Cited Document	Notes
/DB/	3-1	CA 2202572	10/14/1998	Nortel Networks Ltd.	
/DB/	3-2	EP-0824236 A2	02/18/1998	Baird	
/DB/	3-3	EP-0865180 A2	09/16/1998	Lucent	
	3-4				
	3-5				
	3-6				
	3-7				
	3-8				
	3-9				
	3-10				
	3-11				
	3-12				
	3-13				
	3-14				
	3-15				
	3-16				
	3-17				
	3-18				
	3-19				
	3-20				
	3-21				
	3-22				
	3-23				
	3-24				
	3-25				
	3-26				

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	4-1	Awerbuch, B. et al., Distributed Paging for General Networks. In Proc. of the 7th ACM-SIAM Symposium on Discrete Algorithms, pp. 574-583, January 1996.	
/DB/	4-2	Baentsch, M, et al. "Enhancing the Web's Infrastructure: From Caching to Replication." IEEE Internet Computing, 1(2):18--27, Mar. 1997.	
/DB/	4-3	CARTER, J. Lawrence et al., "Universal Classes of Hash Functions", Journal of Computer and System Sciences, Vol. 18, No. 2, (April 1979), 143-154	
/DB/	4-4	Cisco Systems, Inc., "Cache Director System", http://www.cisco.com/warp/public/751/cache/cdsov.htm, (1997), 1-5	
/DB/	4-5	Cisco Systems, Inc., "Cisco Cache Engine", http://www.cisco.com/warp/public/751/cache/index.html	
/DB/	4-6	Gadde, S., et al., "Reduce, reuse, recycle: An approach to building large internet caches," in Workshop on Hot Topics in Operating Systems (HotOS), April 1997.	
/DB/	4-7	Naor, Moni, et al., "The Load, Capacity and Avaliability of Quorum Systems", In Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, (November 1994), 214-225	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	5-1	Pultar, Giray; "USENIX, Automatically Selecting a Close Mirror Based on Network Topology," Proceedings of the Twelfth Systems Administraion Conference; Boston, Massachusetts, Dec. 1998, p. 159-166.	
/DB/	5-2	Ravi, R., "Rapid Rumor Ramification: Approximating the Minimum Broadcast Time", In Proceedings of teh 35th IEEE Symposim on Foundation of Computer Science, (November 1994), 202-213	
/DB/	5-3	Schmidt, J., et al., "Chernoff-Hoeffding Bounds for Applications with Limited Independence", in Proc. 4th ACM-SIAM Symp. on Discrete Algorithms, (1993), 331-340.	
/DB/	5-4	WEGMAN, Mark, et al., "New Hash Functions and Their Use in Authentication and Set Equality", Journal of Computer and System Sciences, Vol. 22, (June 1981), 265 -279	
	5-5		
	5-6		
	5-7		

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	1-1	Korkea-aho, M. (1995). Scalability in Distributed Multimedia Systems, Technical report TKO-B128, Helsinki University of Technology.	
	1-2		
	1-3		
	1-4		
	1-5		
	1-6		
	1-7		

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

OIPE JUN 23 2005 PATENT & TRADEMARK OFFICE JC172

IFW

Docket No.: 6233P001DC3

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:

DAVID A. FARBER, ET AL.

Application No.: 11/065,412

Filed: February 23, 2005

For: **Optimized Network Resource Location**

Art Group: 2157

Examiner: Unassigned

INFORMATION DISCLOSURE STATEMENT UNDER 37 C.F.R. §1.97

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

In accordance with the duty of disclosure, enclosed is a copy of IDS Citation Form PTO/SB/08 or PTO-1449, together with copies of the documents cited on that form, except for copies not required to be submitted (e.g., copies of U.S. patents and U.S. published patent applications need not be enclosed). This IDS and IDS Citation Form are being submitted before the mailing of a first Office Action. It is respectfully requested that the cited references be considered and that the enclosed copy of PTO/SB/08 be initialed by the Examiner to indicate such consideration and a copy thereof returned to applicant(s). Some or all of the references listed on the enclosed PTO/SB/08 were previously identified in the parent application (Application No. 09/612,598, filed July 7, 2000) and copies of the references were furnished at that time. Accordingly, per 37 CFR §1.98(d)(1) additional copies of those references are not submitted herewith.

The submission of this Information Disclosure Statement is not to be construed as a representation that a search has been made in the subject application and is not to be construed as an admission that the information cited in this statement is material to patentability.

Please charge any fees due to Deposit Account 02-2666. A duplicate copy of the Fee Transmittal (PTO/SB/17) is enclosed for this purpose.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: June 21, 2005

William W. Schaal, Reg. No. 39,018

12400 Wilshire Boulevard, 7th Floor
Los Angeles, CA 90025
Telephone: (714) 557-3800

I hereby certify that this correspondence is being deposited with the United States Postal Service on the date shown below with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Susan McFarlane — *06-21-05*

Susan McFarlane — ***Date***

OIPE JC172 JUN 23 2005 PATENT & TRADEMARK OFFICE

TRANSMITTAL FORM

(to be used for all correspondence after initial filing)

Application No.	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Art Unit	2157
Examiner Name	Unassigned
Total Number of Pages in This Submission: 14	Attorney Docket Number: 6233P001DC3

ENCLOSURES ***(check all that apply)***

- ☒ Fee Transmittal Form
  - ☐ Fee Attached
- ☐ Amendment / Response
  - ☐ After Final
  - ☐ Affidavits/declaration(s)
- ☐ Extension of Time Request
- ☐ Express Abandonment Request
- ☒ Information Disclosure Statement
  - ☒ PTO/SB/08
- ☐ Certified Copy of Priority Document(s)
- ☐ Response to Missing Parts/ Incomplete Application
  - ☐ Basic Filing Fee
  - ☐ Declaration/POA
  - ☐ Response to Missing Parts under 37 CFR 1.52 or 1.53
- ☐ Drawing(s)
- ☐ Licensing-related Papers
- ☐ Petition
- ☐ Petition to Convert a Provisional Application
- ☐ Power of Attorney, Revocation Change of Correspondence Address
- ☐ Terminal Disclaimer
- ☐ Request for Refund
- ☐ CD, Number of CD(s)
- ☐ After Allowance Communication to Group
- ☐ Appeal Communication to Board of Appeals and Interferences
- ☐ Appeal Communication to Group **(Appeal Notice, Brief, Reply Brief)**
- ☐ Proprietary Information
- ☐ Status Letter
- ☐ Other Enclosure(s) *(please identify below):*

Remarks

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm or Individual name	William W. Schaal, Reg. No. 39,018 BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
Signature	
Date	June 21, 2005

CERTIFICATE OF MAILING/TRANSMISSION

I hereby certify that this correspondence is being deposited with the United States Postal Service on the date shown below with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Typed or printed name	**Susan McFarlane**		
Signature		Date	**June 21, 2005**

Based on PTO/SB/21 (04-04) as modified by Blakely, Sokoloff, Taylor & Zafman (wlr) 06/04/2004.
SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

OIPE JC172 JUN 23 2005 PATENT & TRADEMARK OFFICE

FEE TRANSMITTAL for FY 2005

Patent fees are subject to annual revision.

☐ Applicant claims small entity status. See 37 CFR 1.27.

TOTAL AMOUNT OF PAYMENT ($) 0.00

Complete if Known	
Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Examiner Name	Unassigned
Art Unit	2157
Attorney Docket No.	6233P001DC3

METHOD OF PAYMENT *(check all that apply)*

☐ Check ☐ Credit card ☐ Money Order ☐ None ☐ Other (please identify): _____

☒ Deposit Account Deposit Account Number: 02-2666 Deposit Account Name: Blakely, Sokoloff, Taylor & Zafman LLP

For the above-identified deposit account, the Director is hereby authorized to: (check all that apply)

☒ Charge fee(s) indicated below ☐ Charge fee(s) indicated below, except for the filing fee

☒ Charge any additional fee(s) or underpayment of fee(s) under 37 CFR §§ 1.16, 1.17, 1.18 and 1.20. ☐ Credit any overpayments

FEE CALCULATION

1. EXTRA CLAIM FEES

			Extra Claims		Fee from below		Fee Paid
Total Claims	4	- 20* =	0	x	50.00	=	$0.00
Independent Claims	2	- 3* =	0	x	200.00	=	$0.00
Multiple Dependent						=	

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description
1202	50	2202	25	Claims in excess of 20
1201	200	2201	100	Independent claims in excess of 3
1203	360	2203	180	Multiple Dependent claim, if not paid
1204	300	2204	150	**Reissue independent claims over original patent
1205	300	2205	150	**Reissue claims in excess of 20 and over original patent

***or number previously paid, if greater, For Reissues, see below*

SUBTOTAL (1) ($) 0.00

2. ADDITIONAL FEES

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description	Fee Paid
1051	130	2051	65	Surcharge - late filing fee or oath	
1052	50	2052	25	Surcharge - late provisional filing fee or cover sheet.	
2053	130	2053	130	Non-English specification	
1251	120	2251	60	Extension for reply within first month	
1252	450	2252	225	Extension for reply within second month	
1253	1,020	2253	510	Extension for reply within third month	
1254	1,590	2254	795	Extension for reply within fourth month	
1255	2,160	2255	1,080	Extension for reply within fifth month	
1401	500	2401	250	Notice of Appeal	
1402	500	2402	250	Filing a brief in support of an appeal	
1403	1,000	2403	500	Request for oral hearing	
1451	1,510	2451	1,510	Petition to institute a public use proceeding	
1460	130	2460	130	Petitions to the Commissioner	
1807	50	1807	50	Processing fee under 37 CFR 1.17(q)	
1806	180	1806	180	Submission of Information Disclosure Stmt	
1809	790	1809	395	Filing a submission after final rejection (37 CFR § 1.129(a))	
1810	790	2810	395	For each additional invention to be examined (37 CFR § 1.129(b))	
Other fee (specify)					

SUBTOTAL (2) ($)

SUBMITTED BY Complete (if applicable)

Name (Print/Type)	William W. Schaal	Registration No. (Attorney/Agent)	39,018	Telephone	(714) 557-3800
Signature				Date	06/21/05

Based on PTO/SB/17 (12-04) as modified by Blakely, Sokoloff, Taylor & Zafman (wlr) 12/15/2004.
SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

OIPE JC172
JUN 23 2005
PATENT & TRADEMARK OFFICE

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 0651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Group Art Unit	2157
Examiner Name	Unknown

Sheet 1 of 9

Attorney Docket No: 6233P001DC3

US PATENT DOCUMENTS

Examiner Initial *	USP Document Number	Publication Date	Name of Patentee or Applicant of cited Document	Class	Subclass	Filing Date If Appropriate
/DB/	US-4,495,570	01/22/1985	Kitajima, Hiroyuki , et al.			01/07/1982
/DB/	US-4,594,704	06/10/1986	Ollivier, Yves			11/16/1984
/DB/	US-4,726,017	02/16/1988	Krum, Harry A., et al.			05/21/1985
/DB/	US-4,839,798	06/13/1989	Eguchi, Hiroyoshi , et al.			11/06/1985
/DB/	US-4,920,432	04/24/1990	Eggers, Derek C., et al.			01/12/1998
/DB/	US-4,922,417	05/01/1990	Churm, Brian R., et al.			10/24/1986
/DB/	US-4,949,187	08/14/1990	Cohen, Jason M.			12/16/1988
/DB/	US-4,949,248	08/14/1990	Caro, Marshall A.			07/15/1988
/DB/	US-5,826,031	10/20/1998	Nielsen, Jakob			06/10/1996
/DB/	US-5,172,413	12/15/1992	Bradley, Graham C., et al.			12/20/1990
/DB/	US-5,253,341	10/12/1993	Rozmanith, Anthony I., et al.			04/11/1991
/DB/	US-5,287,499	02/15/1994	Nemes, Richard M.			05/16/1991
/DB/	US-5,287,537	02/15/1994	Newmark, Rona J., et al.			09/21/1992
/DB/	US-5,291,554	03/01/1994	Morales,			08/19/1992
/DB/	US-5,341,477	08/23/1994	Pitken, Richard P., et al.			08/06/1993
/DB/	US-5,371,532	12/06/1994	Gleman, Alexander , et al.			05/15/1992
/DB/	US-5,410,343	04/25/1995	Coddington, Carl D., et al.			04/02/1993
/DB/	US-5,414,455	05/09/1995	Hooper, Donald F., et al.			07/07/1993
/DB/	US-5,442,389	08/15/1995	Blahut, Donald E., et al.			12/28/1992
/DB/	US-5,442,390	08/15/1995	Hooper, Donald F., et al.			07/07/1993
/DB/	US-5,442,749	08/15/1995	Northcutt, J D., et al.			07/09/1993
/DB/	US-5,471,622	11/28/1995	Eadline, Douglas J.			04/20/1995
/DB/	US-5,475,615	12/12/1995	Lin, K D.			12/23/1993
/DB/	US-5,508,732	04/16/1996	Bottomley, John F., et al.			03/22/1993
/DB/	US-5,515,511	05/07/1996	Nguyen, Tam M., et al.			06/06/1994
/DB/	US-5,519,435	05/21/1996	Anderson, Michael H.			09/01/1994

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	11/065,412
	Filing Date	February 23, 2005
	First Named Inventor	David A. Farber
	Group Art Unit	2157
	Examiner Name	Unknown
Sheet 2 of 9	Attorney Docket No: 6233P001DC3	

/DB/	US-5,528,281	06/18/1996	Grady, John , et al.			07/30/1993
/DB/	US-5,539,621	07/23/1996	Kikinis, Dan			06/23/1995
/DB/	US-5,542,087	07/30/1996	Neimat, Marie-Anne , et al.			10/15/1993
/DB/	US-5,544,313	08/06/1996	Shachnai, Hadas , et al.		05/11/1994	05/11/1994
/DB/	US-5,544,327	08/06/1996	Dan, Asit , et al.			07/26/1994
/DB/	US-5,550,577	08/27/1996	Verbiest, Willem J., et al.			05/19/1994
/DB/	US-5,550,863	08/27/1996	Yurt, Paul , et al.			10/08/1993
/DB/	US-5,550,982	08/27/1996	Long, James E., et al.			06/24/1993
/DB/	US-5,557,317	09/17/1996	Nishio, Makoto , et al.			05/22/1995
/DB/	US-5,572,643	11/05/1996	Judson, David H.			10/19/1995
/DB/	US-5,590,288	12/31/1996	Castor, Patrick F., et al.			03/13/1995
/DB/	US-5,592,611	01/04/1997	Midgely, Christopher W., et al.			03/14/1995
/DB/	US-5,594,910	01/14/1997	Filepp, Robert , et al.			11/26/1993
/DB/	US-5,603,026	02/11/1997	Demers, Alan J., et al.			12/07/1994
/DB/	US-5,619,648	04/08/1997	Canale,			
/DB/	US-5,623,656	04/22/1997	Lyons, Kenneth B.			12/15/1994
/DB/	US-5,625,781	04/29/1997	Cline, Troy L., et al.			10/31/1995
/DB/	US-5,627,829	05/06/1997	Gleeson, Bryan J., et al.			06/06/1995
/DB/	US-5,630,067	05/13/1997	Kindell, Craig N., et al.			06/04/1996
/DB/	US-5,633,999	05/27/1997	Clowes, Richard F., et al.			12/18/1995
/DB/	US-5,634,006	05/27/1997	Baugher, Mark J., et al.			10/11/1995
/DB/	US-5,638,443	06/10/1997	Stefik, Mark J., et al.			11/23/1994
/DB/	US-5,644,714	07/01/1997	Kikinis, Dan			04/10/1995
/DB/	US-5,646,676	05/30/1995	Dewkett, Thomas J., et al.			05/30/1995
/DB/	US-5,649,186	07/15/1997	Ferguson,			08/07/1995
/DB/	US-5,659,729	08/19/1997	Nielsen, Jakob			02/01/1996
/DB/	US-5,666,362	09/09/1997	Chen, Cheng T., et al.			07/25/1995
/DB/	US-5,671,279	09/23/1997	Elgamal, Taher			11/13/1995
/DB/	US-5,682,512	10/28/1997	Tetrick, R. S.			06/30/1995
/DB/	US-5,699,513	12/16/1997	Feigen, Ronald G., et al.	395		03/31/1995

EXAMINER	/Douglas Blair/	**DATE CONSIDERED**	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
• EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Group Art Unit	2157
Examiner Name	Unknown

Sheet 3 of 9

Attorney Docket No: 6233P001DC3

/DB/	US-5,712,979	01/27/1998	Graber, Terry E., et al.			09/20/1995
/DB/	US-5,715,453	02/03/1998	Stuart, Gordon G., et al.			05/31/1996
/DB/	US-5,721,914	02/24/1998	DeVries, Larry R.			09/14/1995
/DB/	US-5,734,831	03/31/1998	Sanders, James B.			04/26/1996
/DB/	US-5,740,423	04/14/1998	Logan, James R., et al.			04/14/1998
/DB/	US-5,742,762	04/21/1998	Scholl, Thomas H., et al.			05/19/1995
/DB/	US-5,751,961	05/12/1998	Smyk, Darek A.			01/31/1996
/DB/	US-5,761,507	06/02/1998	Govett, Ian R.			03/05/1996
/DB/	US-5,761,663	06/02/1998	Lagarde, Konrad C., et al.			05/15/1997
/DB/	US-5,774,660	06/30/1998	Brendel, Juergen , et al.			08/05/1996
/DB/	US-5,774,668	06/30/1998	Choquier, Philippe , et al.			06/07/1995
/DB/	US-5,777,989	07/07/1998	McGarvey, John R.			12/19/1995
/DB/	US-5,784,058	07/21/1998	LaStrange, Thomas E., et al.			05/28/1996
/DB/	US-5,796,952	08/18/1998	Davis, Owen , et al.			03/21/1997
/DB/	US-5,799,141	08/25/1998	Galipeau, Kenneth J., et al.			06/09/1995
/DB/	US-5,802,291	09/01/1998	Balick, Maurice , et al.			03/30/1995
/DB/	US-5,812,769	09/22/1998	Graber, Terry E., et al.			09/20/1995
/DB/	US-5,828,847	10/27/1998	Gehr, Chuck R., et al.			04/19/1996
/DB/	US-5,832,506	11/03/1998	Kuzma, Andrew J.			03/29/1996
/DB/	US-5,835,718	11/10/1998	Blewett, Charles D.			04/10/1996
/DB/	US-5,838,906	11/17/1998	Doyle, Michael D., et al.			10/17/1994
/DB/	US-5,856,974	01/05/1999	Gervais, Joseph L., et al.			01/05/1999
/DB/	US-5,862,339	01/19/1999	Bonnaure, Patrick P.			07/09/1996
/DB/	US-5,867,706	02/02/1999	Martin, Sean J., et al.			02/02/1999
/DB/	US-5,867,799	02/02/1999	Lang, Andrew K., et al.			04/04/1996
/DB/	US-5,870,546	02/09/1999	Kirsch, Steven T.			12/23/1997
/DB/	US-5,870,559	02/09/1999	Leshem, Eran , et al.			04/11/1997
/DB/	US-5,878,212	03/02/1999	Civanlar, Seyhan , et al.			07/31/1995
/DB/	US-5,884,038	03/16/1999	Kapoor, Rupesh			05/02/1997

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Group Art Unit	2157
Examiner Name	Unknown

Sheet 4 of 9

Attorney Docket No: 6233P001DC3

/DB/	US-5,890,171	03/30/1999	Blumer, Thomas P., et al.			08/06/1996
/DB/	US-5,896,533	04/20/1999	Ramos, Daniel , et al.			07/06/1995
/DB/	US-5,903,723	05/11/1999	Beck, Robert D., et al.			12/21/1995
/DB/	US-5,907,704	05/25/1999	Gudmundson, Norman K., et al.			10/02/1996
/DB/	US-5,913,028	06/15/1999	Wang, Frank , et al.			10/06/1995
/DB/	US-5,913,033	06/15/1999	Grout, Gordon B.			12/20/1996
/DB/	US-5,918,010	06/29/1999	Appleman, Kenneth H., et al.			02/06/1998
/DB/	US-5,919,247	07/06/1999	Van Hoff, Arthur , et al.			07/24/1996
/DB/	US-5,933,832	08/03/1999	Suzuoka, Takashi , et al.			09/10/1997
/DB/	US-5,945,989	08/31/1999	Freishtat, Gregg S., et al.			03/25/1997
/DB/	US-5,956,716	09/21/1999	Kenner, Brian , et al.			06/07/1996
/DB/	US-5,958,008	09/28/1999	Pogrebisky, Michael , et al.			04/11/1997
/DB/	US-5,961,596	10/05/1999	Takubo, Shunji , et al.			02/14/1997
/DB/	US-5,968,121	10/19/1999	Logan, Keith W., et al.			08/13/1997
/DB/	US-5,983,214	11/09/1999	Lang, Andrew K., et al.			11/05/1998
/DB/	US-5,983,227	11/09/1999	Nazem, Farzad , et al.			06/12/1997
/DB/	US-5,991,809	11/23/1999	Kriegsman, Mark E.			07/25/1997
/DB/	US-6,006,264	12/21/1999	Colby, Steven , et al.			03/30/1998
/DB/	US-6,012,090	01/04/2000	Chung, Pi-yu , et al.	709		03/14/1997
/DB/	US-6,014,698	01/11/2000	Griffiths, Michael J.			06/11/1997
/DB/	US-6,029,175	02/22/2000	Chow, Yen-whei , et al.			06/07/1996
/DB/	US-6,035,332	03/07/2000	Ingrassia, Jr., Michael I., et al.	709		10/06/1997
/DB/	US-6,038,310	03/14/2000	Hollywood, Steven , et al.			08/01/1995
/DB/	US-6,041,324	03/21/2000	Earl, Joel R., et al.			11/17/1997
/DB/	US-6,052,718	04/18/2000	Gifford, David K.			01/07/1997
/DB/	US-6,052,730	04/18/2000	Feliciano, Ramon M., et al.			01/09/1998
/DB/	US-6,065,051	05/16/2000	Steele, Douglas W., et al.			04/15/1998

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Group Art Unit	2157
Examiner Name	Unknown

Sheet 5 of 9

Attorney Docket No: 6233P001DC3

/DB/	US-6,065,062	05/16/2000	Periasamy, Ravi , et al.			12/10/1997
/DB/	US-6,070,191	05/30/2000	Narendran, Balakrishnan , et al.			10/17/1997
/DB/	US-6,081,829	06/27/2000	Sidana, Ashmeet S.			01/31/1996
/DB/	US-6,092,112	07/18/2000	Fukushige, Yoshio			01/14/1997
/DB/	US-6,105,028	08/15/2000	Sullivan, David J., et al.			06/26/1997
/DB/	US-6,108,673	08/22/2000	Brandt, Marcia L., et al.			02/25/1997
/DB/	US-6,108,703	08/22/2000	Leighton, F. T., et al.			05/19/1999
/DB/	US-6,112,231	08/29/2000	DeSimone, Antonio , et al.			10/27/1998
/DB/	US-6,112,240	08/29/2000	Pogue, Michael A., et al.			09/03/1997
/DB/	US-6,128,601	10/03/2000	Van Horne, Peter , et al.			03/02/1998
/DB/	US-6,128,660	10/03/2000	Grimm, Steven M., et al.			03/20/1997
/DB/	US-6,144,375	11/07/2000	Jain, Ramesh C., et al.			08/14/1998
/DB/	US-6,163,779	12/19/2000	Mantha, Murthy L., et al.			09/29/1997
/DB/	US-6,173,322	01/09/2001	Hu, Wei-Ming			06/05/1997
/DB/	US-6,178,160	01/23/2001	Bolton, Derek W., et al.			12/23/1997
/DB/	US-6,185,598	02/06/2001	Farber, David A.			02/10/1998
/DB/	US-6,185,619	02/06/2001	Joffe, Rodney L., et al.	B1		11/07/1997
/DB/	US-6,189,030	02/13/2001	Kirsch, Steven T., et al.			05/01/1998
/DB/	US-6,226,642	05/01/2001	Beranek, Michael J., et al.			09/11/1997
/DB/	US-6,256,675	07/03/2001	Rabinovich, Michael			05/06/1998
/DB/	US-6,286,045	09/04/2001	Griffiths, Michael J., et al.			05/19/1997

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Foreign Document No	Publication Date	Name of Patentee or Applicant of cited Document	Class	Subclass	T[2]
/DB/	CA-2202572	10/14/1998	Law, Ka LUn E., et al.			
/DB/	EP-0801487A2	10/15/1997	Charles, Douglas B.			
/DB/	EP-0824236A2	02/18/1998	Baird, Richard , et al.			
/DB/	EP-0865180A2	09/16/1998	Chung, Pi-Yu , et al.			

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Complete if Known

Application Number	11/065,412
Filing Date	February 23, 2005
First Named Inventor	David A. Farber
Group Art Unit	2157
Examiner Name	Unknown

Sheet 6 of 9

Attorney Docket No: 6233P001DC3

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Foreign Document No	Publication Date	Name of Patentee or Applicant of cited Document	Class	Subclass	T[2]
/DB/	GB-2281793	03/15/1995	Wallis, Graham D., et al.			
/DB/	WO-96/42041	12/27/1996	Levergood, Thomas M., et al.			
/DB/	WO-97/11429	03/27/1997	Graber, Terry E., et al.			
/DB/	WO-97/29423	08/14/1997	Stanford-Clark, Andrew , et al.			
/DB/	WO-98/04985	02/05/1998	Kriegsman, Mark			
/DB/	WO-98/06033	02/12/1998	Agranat, Ian D., et al.			

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/DB/		"Cache Director System", http:/www.cisco.com/warp/public/751/cache/cds ov.htm, Cisco Systems, Inc.,(1997),1-5	
/DB/		"Cache Director Systems Technology Overview", http:/www.cisco.com/warp/public/751/cache/cds ds.htm, Cisco System, Inc.,(1997),1-3	
/DB/		"Cisco Cache Engine", http:/www.cisco.com/warp/public/751/cache/index.html, Cisco Systems, Inc.,	
/DB/		"Cisco Introduces Web Cache Product for Scaling the Internet", http:/www.cisco.com/warp/public;146/1947.html, Cisco Systmes, Inc.,(1997),	
/DB/		"Cisco Systmes Ships Cisco LocalDirector to Meet Demands of High-volume Web Traffic", http://www.cisco.com/wrap/public/146/323.html, Cisco Systems, Inc.,(1996),1-2	
/DB/		"Cisco Take Global Route", PC Week News, (2/17/1997),23	
/DB/		"Communication (1 page) and European Search Report for", Application No. EP 00 12 8346 (3 pages), (08/24/2001),	
/DB/		"Excite Chooses Cisco Local Director to Support its Growing Service", http:/www.cisco.com/wrap/public/146/1857.html, Cisco Systems, Inc.,(1997),1-2	
/DB/		"Global IP/PX Service Should Keep Network Delays Down", Infoworld, (01/20/1997),1 page	
/DB/		"How to Cost-Effectively Scale Web Servers", http://www.cisco.com/warp/public/784//5.html, Cisco Systems, Inc.,(1996),1-5	
/DB/		"http://www.cisco.com/warp/public/751/distdir/dd wp.htm", Cisco Systems Inc., (1997),1-15	
/DB/		"http://www.cisco.com/wrap/public/751/lodir/lodir ds.htm", Cisco System Inc., (1997),1-4	
/DB/		"International Serach Report for PCT/US99/01477", (06/08/1999),	
/DB/		"Local Area Network Server Replacement Procedure", IBM Technical Disclosure Bulletin, Vol. 38, No. 1,(January 1995),235-236	
/DB/		"Local Director", http://.cisco.com/warp/public/751/lodir/index.html, 2 pages	

EXAMINER /Douglas Blair/ **DATE CONSIDERED** 09/15/2008

Substitute Disclosure Statement Form (PTO-1449)

* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	11/065,412
	Filing Date	February 23, 2005
	First Named Inventor	David A. Farber
	Group Art Unit	2157
	Examiner Name	Unknown
Sheet 7 of 9	Attorney Docket No: 6233P001DC3	

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS			
Examiner Initials*	**Cite No [1]**	**Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.**	**T[2]**
/DB/		"Local Director", Cisco Systems, Inc., 1-3	
/DB/		"Patent Abstracts of Japan, Electronic Mail Multiplexing System and Communication Control Method in the System", JP 05162529, (06/30/1993),	
/DB/		"Patent Abstracts of Japan, Method and Device for Repeating and Converting Information", JP 08328583, (09/1996),	
/DB/		"Scaling the Internet News Service", http:/www.cisco.com/warp/public/751/lodir/news wp.htm, Cisco Systems, Inc.,(1997),1-11	
/DB/		"Scaling the World Wide Web", http://www.cisco.com/warp/public/751/lodir/swww wp.html, Cisco Systems, Inc.,(1996),1-7	
/DB/		ADLER, "Distributed Coordination Models for Client/Server Computing", Comptuer, Vol. 28, No. 4,(04/01/1995),14-22	
/DB/		ANDRESEN, et al., "SWEB: Towards a Scalable World Wide Web Server on Multicomputers", Proceedings of IPPS, (04/15/1996),850-856	
/DB/		BESTAVROS, AZER , "Speculative Data Dissermination and Service to Reduce Server Load, Network Traffic and Service Time in Distributed Information Systems", In Proceedings of ICDE '96: The 1996 International Conference Conference on Data Engineering, (March 1996),4 pages	
/DB/		CARTER, J. LAWRENCE , et al., "Universal Classes of Hash Functions", Journal of Computer and System Sciences, Vol. 18, No. 2, (April 1979),143-154	
/DB/		CHANKHUNTHOD, ANAWAT , et al., "A Hierarchical Internet Object Cache", USENIX Proceedings, (January 1996),153-163	
/DB/		CORMEN, THOMAS H., et al., "Introduction to Algorithms", The MIT Press, Cambridge, Massachusetts, (1994),219-243, 991-993	
/DB/		DE BRA, P.M.E. , et al., "Information Retrieval inthe World Wide Web: Making client-based searching feasible", Computer Networks and ISDN Systems, NL, North Holland Publishing, Amsterdam, Vol. 27. No. 2, ISSN: 0169-7552, (11/01/1994),183-192	
/DB/		DEERING, STEPHEN , et al., "Multicast Routing in Datagram Internetworks and Extended LANs", ACM Transactions on Computer Systems, Vol. 8, No. 2, (May 1990),85-110	
/DB/		DEVINE, ROBERT , "Design and Implementation of DDH: A Distributed Dynamic Hashing Algorithm", In Proceedings of 4th International Conference on Foundations of Data Organizations and Algorithms, (1993),101-114	
/DB/		FEELEY, MICHAEL , et al., "Implementing Global Memory Management in a Workstation Cluster", In Proceedings of the 15th ACM Symposium on Operating Systems Principles, (1995),201-212	
/DB/		FLOYD, SALLEY , et al., "A Reliable Multicast Framework for Light-Weight Sessions and Application Level Framing", In Proceeding of ACM SIGCOMM '95, 342-356	

EXAMINER	/Douglas Blair/	**DATE CONSIDERED**	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional) [2] Applicant is to place a check mark here if English language Translation is attached

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	11/065,412
	Filing Date	February 23, 2005
	First Named Inventor	David A. Farber
	Group Art Unit	2157
	Examiner Name	Unknown
Sheet 8 of 9	Attorney Docket No: 6233P001DC3	

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS			
Examiner Initials*	**Cite No [1]**	**Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.**	**T[2]**
/DB/		FREDMAN, MICHAEL , et al., "Storing a Sparse Table with 0(1) Worst Case Access Time", Journal of the Association for Computing Machinery, Vol. 31, No. 3, (July 1984),538-544	
/DB/		GRIGNI, MICHELANGELO , et al., "Tight Bounds on Minimum Broadcasts Networks", SIAM Journal of Discrete Mathematics, Vol. 4, No. 2, (May 1991),207-222	
/DB/		GWERTZMAN, JAMES , et al., "The Case for Geographical Push-Caching", Technical Report HU TR 34-94 (excerpt), Harvard University, DAS, Cambridge, MA 02138, (1994),2 pages	
/DB/		GWERTZMAN, J.S. , et al., "The case for Geographical push-caching", Workshop on Hot Topics in Operating Systems, (05/04/1995),51-55	
/DB/		GWERTZMAN, JAMES , et al., "World-Wide Web Cache Consistency", In Proceedings of the 1996 USENIX Technical Conference, (January 1996),8 pages	
/DB/		KARGER, DAVID , et al., "Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web", In Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Coomputing, (May 1997),654-663	
/DB/		LITWIN, WLTHOLD , et al., "LH-A Scaleable, Distributed Data Structure", ACM Transactions on Database Systems, Vol. 21, No. 4, (December 1996),480-525	
/DB/		MALPANI, RADHIKA , et al., "Making World Wide Web Caching Servers Cooperate", In Proceedings of World Wide Web Conference, (1996),6 pages	
/DB/		MOURAD, et al., "Scalable Web Server Architectures", Proceedings IEEE Symposium on Computers and Communications, (07/01/1997),12-16	
/DB/		NAOR, MONI , et al., "The Load, Capacity and Availability of Quorum Systems", In Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, (November 1994),214-225	
/DB/		NISAN, NOAM , "Psuedorandom Generations for Sapce-Bounded Computation", In Proceedings of the Twenty-Second Annual ACM Symposium on Therory of Computing, (May 1990),204-212	
/DB/		PALMER, MARK , et al., "Fido: A Cache that Learns to Fetch", In Proceedings of the 17th International Conference on Very Large Data Bases, (September 1991),255-264	
/DB/		PANIGRAHY, RINA , "Relieving Hot Spots onthe World Wide Web", Massachusetts Institute of Technology, (June 1997),1-66	
/DB/		PELEG, DAVID , et al., "The Availability of Quorum Systems", Inforamtion and Computation, 123, (1995),210-223	
/DB/		PLAXTON, C. GREG , et al., "Fast Fault-Tolerant Concurrent Access to Shared Objects", In Proceedings of the 37th IEEE Symposium of Foundations of Computer Science, (1996),570-579	
/DB/		RABIN, MICHAEL , "Efficient Dispersal of Information for Security, Load	

EXAMINER	/Douglas Blair/	**DATE CONSIDERED**	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. 1 Applicant's unique citation designation number (optional) 2 Applicant is to place a check mark here if English language Translation is attached

PTO/SB/08A(10-01)
Approved for use through 10/31/2002. OMB 651-0031
US Patent & Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449A/PTO **INFORMATION DISCLOSURE STATEMENT BY APPLICANT** *(Use as many sheets as necessary)*	*Complete if Known*	
	Application Number	11/065,412
	Filing Date	February 23, 2005
	First Named Inventor	David A. Farber
	Group Art Unit	2157
	Examiner Name	Unknown
Sheet 9 of 9	Attorney Docket No: 6233P001DC3	

OTHER DOCUMENTS -- NON PATENT LITERATURE DOCUMENTS			
Examiner Initials*	**Cite No [1]**	**Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.**	**T[2]**
/DB/		Balancing and Fault Tolerance", Journal of the ACM, Vol. 36, No. 2, (April 1989),335-348	
/DB/		RAVI, R. , "Rapid Rumor Ramification: Approximating the Minimum Broadcast Time", In Proceedings of the 35th IEEE Symposim on Foundation of Computer Science, (November 1994),202-213	
/DB/		SCHMIDT, JEANETTE , et al., "Chernoff-Hoeffding Bounds for Applications with Limited Independence", In Proceedings of the 4th ACS-SIAM Symposium on Discrete Algorithms, (1993),331-340	
/DB/		TARJAN, ROBERT E., et al., "Storing a Sparse Table", Communications of the ACM, Vol. 22, No. 11, (November 1979),606-611	
/DB/		VITTER, JEFFREY S., et al., "Optimal Prefetching via Data Compression", In Proceedings of 32nd IEEE Symposium on Foundations of Computer Science, (November 1991),121-130	
/DB/		WEGMAN, MARK , et al., "New Hash Functions and Their Use in Authentication and Set Equality", Journal of Computer and System Sciences, Vol. 22, (June 1981),265-279	
/DB/		YAO, ANDREW C., "Should Tables be Sorted?", Journal of Association for Computing Machinery, Vol. 28, No. 3, (July 1981),615-628	

EXAMINER	/Douglas Blair/	**DATE CONSIDERED**	09/15/2008

Substitute Disclosure Statement Form (PTO-1449)
* EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional) [2] Applicant is to place a check mark here if English language Translation is attached

OIPE 1AP44
JUL 12 2007
PATENT & TRADEMARK OFFICE

DFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2157
Application Serial No.: 11/065,412	Examiner: Unassigned
Application Filing Date: February 23, 2005	Date: July 12, 2007
Title: **OPTIMIZED NETWORK RESOURCE LOCATION**	Confirmation No.: 3296

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached (except for copies not required to be submitted).

Copies of some the references were submitted to, or cited by, the Office in parent Application No. 09/612,598, filed July 7, 2000, and/or grandparent Application No. 09/021,506, filed February 10, 1998, which are relied upon for an earlier effective filing date under 35 U.S.C. 120, and copies of the references were furnished at that time. Accordingly, copies of those references are not submitted herewith. 37 C.F.R. §1.98(d).

It is respectfully requested that the cited references be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom. It is further respectfully requested that the enclosed form(s) PTO-1449

be initialed by the Examiner to indicate such consideration, and that a copy thereof be returned to applicant.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed within three (3) months of the U.S. filing date OR before the mailing date of a first Office Action on the merits. No certification or fee is required.

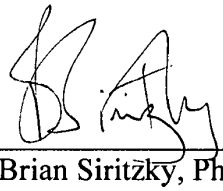
RELATED APPLICATIONS

The Examiner's attention is directed to the following related applications:

09/021,506 (filed 02-10-1998, now U.S. Patent No. 6,185,598);
09/612,598 (filed 07-07-2000, abandoned);
09/930,955 (filed 08-17-2001, abandoned);
10/004,430 (filed 11-25-2003, now U.S. Patent No. 6,554,807);
10/095,811 (filed 03-13-2002, now U.S. Patent No. 7,054,935);
11/441,253 (filed 05-26-2006, Pending);
11/806,147 (filed 05-30-2007, Pending);
11/806,152 (filed 05-30-2007, Pending);
11/806,153 (filed 05-30-2007, Pending); and
11/806,154 (filed 05-30-2007, Pending).

All of the above applications are directed to related technical subject matter. The identification of these U.S. Patent Applications is not to be construed as a waiver of secrecy as to those applications now or upon issuance of the present application as a patent. The Examiner is respectfully requested to consider art cited therein during examination.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 42624	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

PTO/SB/08a (07-05)
Approved for use through 06/30/2006. OMB 0651-0031
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
JUL 12 2007
PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
FORM PTO-1449 (modified)

Sheet 1 of 9

Application No.	11/065,412
Filing Date	Feb. 23, 2005
First Named Inventor	Farber, David A.
Group Art Unit	2157
Examiner Name	unassigned
Attorney Docket No.	2711-0073
Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	1-1	US-2004/0139097	15-Jul-04	Farber et al.
/DB/	1-2	US-2005/0038851	17-Feb-05	Kriegsman
/DB/	1-3	US-2005/0114296	26-May-05	Farber et al.
/DB/	1-4	US-4591983	27-May-86	Bennett et al.
/DB/	1-5	US-4594704	Jun-86	Ollivier
/DB/	1-6	US-4803641	7-Feb-89	Hardy et al.
/DB/	1-7	US-4847784	11-Jul-89	Clancey
/DB/	1-8	US-4943932	24-Jul-90	Lark et al.
/DB/	1-9	US-5029232	02-Jul-91	Nall
/DB/	1-10	US-5130792	14-Jul-92	Tindell et al.
/DB/	1-11	US-5132992	21-Jul-92	Yurt et al.
/DB/	1-12	US-5136716	08/1992	Churm et al.
/DB/	1-13	US-5191573	02-Mar-93	Hair
/DB/	1-14	US-5253275	12-Oct-93	Yurt et al.
/DB/	1-15	US-5675734	7-Oct-97	Hair
/DB/	1-16	US-5764906	06/1998	Edelstein et al.
/DB/	1-17	US-5802106	01-Sep-98	Packer
/DB/	1-18	US-5815664	09/1998	Asano
/DB/	1-19	US-5832514	11/1998	Norin et al.
/DB/	1-20	US-5845303	01-Dec-98	Templeman
/DB/	1-21	US-5893116	04/1999	Simmonds et al.
/DB/	1-22	US-5894554	05/1999	Lowery et al.
/DB/	1-23	US-5896533	04/1999	Ramos et al.
/DB/	1-24	US-5920701	07/1999	Miller et al.
/DB/	1-25	US-5935207	08/1999	Logue et al.
/DB/	1-26	US-5956489	09/1999	San Andres et al.

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	2-1	US-5966440	12-Oct-99	Hair
/DB/	2-2	US-5978791	02-Nov-99	Farber et al.
/DB/	2-3	US-5987606	11/1999	Cirasole et al.
/DB/	2-4	US-6002720	14-Dec-99	Yurt et al.
/DB/	2-5	US-6003030	14-Dec-99	Kenner et al.
/DB/	2-6	US-6014686	01/2000	Elnozahy et al.
/DB/	2-7	US-6018516	25-Jan-00	Packer
/DB/	2-8	US-6026440	02/2000	Sharder et al.
/DB/	2-9	US-6029176	22-Feb-00	Cannon
/DB/	2-10	US-6038216	14-Mar-00	Packer
/DB/	2-11	US-6038610	03/2000	Belfiore et al.
/DB/	2-12	US-6044405	28-Mar-00	Driscoll, III, et al.
/DB/	2-13	US-6046980	04-Apr-00	Packer
/DB/	2-14	US-6049831	11-Apr-00	Gardell et al.
/DB/	2-15	US-6065051	05/2000	Steele et al.
/DB/	2-16	US-6092204	07/2000	Baker
/DB/	2-17	US-6112239	29-Aug-00	Kenner et al.
/DB/	2-18	US-6115357	05-Sep-00	Packer et al.
/DB/	2-19	US-6119143	09/2000	Dias et al.
/DB/	2-20	US-6125388	26-Sep-00	Reisman
/DB/	2-21	US-6134583	10/2000	Herriot
/DB/	2-22	US-6144702	07-Nov-00	Yurt et al.
/DB/	2-23	US-6151624	11/2000	Teare et al.
/DB/	2-24	US-6154753	28-Nov-00	McFarland
/DB/	2-25	US-6167427	12/2000	Rabinovich et al.
/DB/	2-26	US-6173311	01/2001	Hassett et al.

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS				
Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	3-1	US-6181867	30-Jan-01	Kenner et al.
/DB/	3-2	US-6230196	05/2001	Guenthner et al.
/DB/	3-3	US-6243760	06/2001	Armbruster et al.
/DB/	3-4	US-6266699	07/2001	Sevcik
/DB/	3-5	US-6269394	22-Apr-03	Kenner et al.
/DB/	3-6	US-6282569	08/2001	Wallis et al.
/DB/	3-7	US-6298041	02-Oct-01	Packer
/DB/	3-8	US-6314565	06-Nov-01	Kenner et al.
/DB/	3-9	US-6347085	02/2002	Kelly
/DB/	3-10	US-6360256	03/2002	Lim
/DB/	3-11	US-6370571	04/2002	Medin, Jr.
/DB/	3-12	US-6370580	09-Apr-02	Kriegsman
/DB/	3-13	US-6412000	25-Jun-02	Riddle et al.
/DB/	3-14	US-6415280	02-Jul-02	Farber et al.
/DB/	3-15	US-6421726	16-Jul-02	Kenner et al.
/DB/	3-16	US-6456630	24-Sep-02	Packer et al.
/DB/	3-17	US-6480893	12-Nov-02	Kriegsman
/DB/	3-18	US-6484204	11/2002	Rabinovich
/DB/	3-19	US-6496856	17-Dec-02	Kenner et al.
/DB/	3-20	US-6502125	31-Dec-02	Kenner et al.
/DB/	3-21	US-6502215	12/2002	Raad et al.
/DB/	3-22	US-6557054	29-Apr-03	Reisman
/DB/	3-23	US-6581090	17-Jun-03	Lindbo et al.
/DB/	3-24	US-6591299	08-Jul-03	Riddle et al.
/DB/	3-25	US-6611862	26-Aug-03	Reisman
/DB/	3-26	US-6658464	02-Dec-03	Reisman

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS				
Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
/DB/	4-1	US-6665706	16-Dec-03	Kenner et al.
/DB/	4-2	US-6741563	25-May-04	Packer
/DB/	4-3	US-6799221	28-Sep-04	Kenner et al.
/DB/	4-4	US-6915329	05-Jul-05	Kriegsman
/DB/	4-5	US-6928442	09-Aug-05	Farber et al.
/DB/	4-6	US-7047300	05/2006	Oehrke et al.
	4-7			
	4-8			
	4-9			
	4-10			
	4-11			
	4-12			
	4-13			
	4-14			
	4-15			
	4-16			
	4-17			
	4-18			
	4-19			
	4-20			
	4-21			
	4-22			
	4-23			
	4-24			
	4-25			
	4-26			

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication Date	Name of Patentee or Applicant of Cited Document	T[1]
/DB/	5-1	EP-0801487 A2	10/15/1997	Law, et al.	
/DB/	5-2	EP-0817444	01/1998	Sun Microsystems	
/DB/	5-3	JP 10-027148	27/01/1998	Shuji, Ono et al.	
/DB/	5-4	JP-05162529	06/1993	Hiroshi, Saito	
/DB/	5-5	JP-08328583	09/1996	NT&T	
/DB/	5-6	JP7066829	03/1995	Hiroshi, Saito	
	5-7				
	5-8				
	5-9				
	5-10				
	5-11				
	5-12				
	5-13				
	5-14				
	5-15				
	5-16				
	5-17				
	5-18				
	5-19				
	5-20				
	5-21				
	5-22				
	5-23				
	5-24				
	5-25				
	5-26				

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 6 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
/DB/	6-1	"Exporting Web Server Final Report," http://www.cs.technion.ac.il/Labs/Lccn/projects/spring97/project4/final_report.html, Spring 1997 (downloaded 07/07/07).	
/DB/	6-2	Andresen, D., et al., Multiprocessor scheduling with client resources to improve the response time of WWW applications, Proc. 11th Int'l Conf. on Supercomputing (Austria, July, 1997). ICS '97. ACM Press, NY, NY, 92-99.	
/DB/	6-3	Basturk, E., et al., "Using network layer anycast for load distribution in the Internet," Tech. Rep., IBM T.J. Watson Research Center, 1997 (21 pgs.).	
/DB/	6-4	Berners-Lee et al., RFC 1738--Uniform Resource Locators, Dec. 1994.	
/DB/	6-5	Bestavros, et al., "Server-Initiated Document Dissemination for the WWW," IEEE Data Engineering Bulletin 19(3):3-11, Sept. 1996.	
/DB/	6-6	Bhattacharjee et al., "Application-layer anycasting," in Proc. IEEE INFOCOM '97, 1997.	
/DB/	6-7	Braun, H., et al., "Web traffic characterization: an assessment of the impact of caching documents from NCSA's web server," Comput. Netw. ISDN Syst. 28, 1-2 (Dec. 1995), 37-51.	
/DB/	6-8	Brisco, T. P. RFC 1794: DNS support for load balancing, Apr. 1995.	
/DB/	6-9	Carter et al., "Dynamic server selection using bandwidth probing in wide-area networks," Tech. Rep. BU-CS-96-007, Comp. Sci. Dept., Boston University, Mar. 1996.	
/DB/	6-10	Carter et al., Server selection using dynamic path characterization in Wide-Area Networks, IEEE INFOCOM '97, 1997 (pp. 1014-1021).	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 7 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
/DB/	7-1	Cisco Systems, Inc., Cisco DistributedDirector 2500 Series Installation and Configuration Guide, pp. xix - xxii; 1-1 to 1-12; 6-1 to 6-18; 7-1 to 7-18; 8-1 to 8-24, pub. date unknown, downloaded 04/07: http://www.cisco.com/univercd/cc/td/doc/product/iaabu/distrdir/dd2501/.	
/DB/	7-2	Cisco Systems, Inc., Cisco DistributedDirector 4700-M Installation and Configuration Guide, pgs. xix-xxii; 1-1 to 1-14; 7-1 to 7-18, 8-1 to 8-20; pub. date unknown, [downloaded 04/07 from http://www.cisco.com/univercd/cc/td/doc/product/iaabu/distrdir/dd4700m/].	
/DB/	7-3	Cohen, J., et al., "Cache Array Routing Protocol v1.1", 29 Sept. 1997; http://tools.ietf.org/id/draft-vinod-carp-v1-01.txt (Last-Modified: Wed, 01 Oct 1997).	
/DB/	7-4	Colajanni, M. and Yu, P. S. 1997. Adaptive TTL schemes for load balancing of distributed Web servers. SIGMETRICS Perform. Eval. Rev. 25, 2 (Sep. 1997), 36-42.	
/DB/	7-5	Crovella et al., Dynamic server selection in the Internet, 3rd IEEE Workshop on the Arch. and Implementation of High Performance Computer Sys. '95, pp. 158-162, Aug. 1995.	
/DB/	7-6	Danzig, P. B., et al., "An analysis of wide-area name server traffic: a study of the Internet Domain Name System," Conf. Proc. Communications Architectures & Protocols (Aug. 1992). D. Oran, Ed. SIGCOMM '92. ACM Press, New York, NY, 281-292.	
/DB/	7-7	Doi, K. "Super Proxy Script--How to make distributed proxy servers by URL hashing," Sharp Corp., http://naragw.sharp.co.jp/sps/, dates unknown (1996 -2000), download 07/07/07.	
/DB/	7-8	Fox, A., "A Framework for Separating Server Scalability and Availability from Internet Application Functionality," PhD thesis, University of California, Berkeley, 1998.	
/DB/	7-9	Fox, A., et al,, "Cluster-based scalable network services", Proc. 16th ACM Symp. on Operating Systems Principles (Saint Malo, France, Oct. 05 - 08, 1997). W. M. Waite, Ed. SOSP '97. ACM Press, New York, NY, 78-91.	
/DB/	7-10	Goldszmidt, et al., "Load Distribution for Scalable Web Servers: Summer Olympics 1996--A Case Study," In Proc. 8th IFIP/IEEE Int'l Workshop on Distributed Systems: Operations and Management, Sydney, Australia, Oct. 1997.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 8 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
/DB/	8-1	Gulbrandsen, A., et al., "A DNS RR for specifying the location of services (DNS SRV)", Network Working Group, RFC 2052, Oct. 1996.	
/DB/	8-2	Guyton et al., "Locating nearby copies of replicated Internet servers," Proc. ACM SIGCOMM '95, pp. 288-298, 1995.	
/DB/	8-3	Jeffrey et al., Proxy-Sharing Proxy Servers, IEEE, pp. 116-119, 1996.	
/DB/	8-4	Kwan et al., NCSA's World Wide Web Server: Design and Performance, IEEE, pp. 68-74, Nov. 1995.	
/DB/	8-5	Luotonen et al., World-Wide Web Proxies, CERN, Apr. 1994 (modified May 24, 1994).	
/DB/	8-6	Mockapetris et al., "Development of the Domain Name System," Proc. SIGCOMM '88 Computer Communications Review, vol. 18, No. 4, Aug. 1988.	
/DB/	8-7	Mockapetris, P., RFC 1034: Domain Names - Concepts and Facilities, Nov. 1987.	
/DB/	8-8	Mockapetris, P., RFC 1035: Domain Names - Implementation and Specification, Nov. 1987.	
/DB/	8-9	Oguchi et al., A Study of Caching Proxy Mechanisms Realized on Wide Area Distributed Networks, High Performance Distributed Computing, 5th Int'l Symposium, pp. 443-449, 1996.	
/DB/	8-10	Peleg D., et al., "The Availability of Quorum Systems," Tech. Rep. CS93-17, The Weizmann Institute of Science, Rehovot, Israel, 1993.	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 9 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
/DB/	9-1	Petri S., et al., "Load Balancing and Fault Tolerance in Workstation Clusters. Migrating Groups of Communicating Processes.", Operating Systems Review, Vol. 29, No. 4, Oct.1995, pp. 25-36.	
/DB/	9-2	Postel, RFC 1591--Domain Name System Structure and Delegation, Mar. 1994.	
/DB/	9-3	Ross, K.W., "Hash-Routing for Collections of Shared Web Caches", IEEE Network Magazine, 11, 7:37-44, Nov-Dec 1997.	
/DB/	9-4	Schemers, R., "lbnamed - A load balancing name server written in Perl," 1995 LISA IX – Sept. 17-22, 1995 – Monterey, CA.	
/DB/	9-5	Schuba, Christoph; "Addressing Weaknesses in the Domain Name System Protocol," COAST Laboratory, Dept. of Computer Sciences, Purdue University; West Layfayette, IN; Aug. 1993, p. 1-87.	
/DB/	9-6	Smith, "What can Archives offer the World Wide Web?", Technical Report 11, University of Kent, Computing Laboratory, University of Kent, Canterbury, UK, 1994.	
/DB/	9-7	Thaler, D. G. and Ravishankar, C. V. 1998. Using name-based mappings to increase hit rates. IEEE/ACM Trans. Netw. 6, 1 (Feb. 1998), 1-14.	
/DB/	9-8	Vixie, Paul; "Name Server Operations Guide for BIND," Internet Software Consortium; La Honda, CA; p. SMM:10-2--SMM:10-30 (undated, 1996).	
/DB/	9-9	Wessels, Intelligent Caching for World-Wide Web Objects, Masters Thesis, University of Colorado, 1995.	
	9-10		

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	1-1	Claims of European Patent No. 1 143 337 B allowed/maintained in oral proceedings, April 24, 2008 [7 pages].	
/DB/	1-2	Communication from European Patent Office dated Feb. 8, 2007 regarding opposition to European Patent No. 1 143 337 B [11 pages].	
/DB/	1-3	Notice of Opposition to a European Patent filed against European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 23, 2004 [5 pages].	
/DB/	1-4	Owner's Response of April 10, 2007 to summons to attend oral proceedings dated February 8, 2007 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 1 and 2 and attachment discussing claim features [26 pages].	
/DB/	1-5	Owner's Response of March 20, 2008 to summons to attend oral proceedings dated January 24, 2008 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 3 to 7 [21 pages].	
/DB/	1-6	Reply of Dec. 21, 2004 of the Patent Proprietor to the Notice of Opposition of European Patent No. 1 143 337 B, Responsive to the Notice of Opposition of Akamai Technologies, Inc. dated April 23, 2004 [11 pages].	
/DB/	1-7	Statement of Opposition to European Patent No. 1 143 337 B, filed by Akamai Technologies, Inc. April 23, 2004 [43 pages].	

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
/DB/	2-1	Written Submissions in Opposition to European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 5, 2007 [5 pages].	
	2-2		
	2-3		
	2-4		
	2-5		
	2-6		
	2-7		

Examiner Signature	/Douglas Blair/	Date Considered	09/15/2008

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11065412	FARBER ET AL.
	Examiner	Art Unit
	DOUGLAS B BLAIR	2142

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	09/15/2008								
	25	✓								

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	11065412	FARBER ET AL.
	Examiner	Art Unit
	DOUGLAS B BLAIR	2142

SEARCHED			
Class	Subclass	Date	Examiner
709	217,219,245	9/15/2008	DBB

SEARCH NOTES		
Search Notes	Date	Examiner
East text search	9/15/2008	DBB

INTERFERENCE SEARCH			
Class	Subclass	Date	Examiner

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	PATENT NUMBER	GROUP ART UNIT	FILE WRAPPER LOCATION
11/065,412		2142	

OC000000032024820

Correspondence Address/Fee Address Change

The following fields have been set to Customer Number 75948 on 09/10/2008

- **Correspondence Address**
- **Maintenance Fee Address**
- **Power of Attorney Address**

The address of record for Customer Number 75948 is:

75948
DAVIDSON BERQUIST JACKSON & GOWDEY, LLP
ATTN: BRIAN SIRITZKY, Ph.D.
4300 WILSON BLVD., 7TH FLOOR
ARLINGTON, VA 22203

PART 1 - ATTORNEY/APPLICANT COPY
page 1 of 1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: August 24, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the references listed on the attached PTO-1449. One copy of each reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

This application is related to U.S. Patent No. 6,185,598 which issued from U.S. application no. 09/021,506 filed February 10, 1998 (hereinafter "the '506 Application"). European Patent No. 1143337 (hereinafter "EP 337"), granted July 23, 2003, also claimed priority to the '506 application. On April 23, 2004, EP 337 was opposed under European Patent Law by Akamai Technologies, Inc. Submitted herewith are documents relating to that opposition. An oral proceeding was held in the European Patent Office (EPO) on April 24, 2008. A copy of the claims allowed/maintained in that oral proceeding is also enclosed.

The following documents are enclosed (and listed on the herewith Form 1449):

- Notice of Opposition to a European Patent filed against European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 23, 2004 [5 pages].
- Statement of Opposition to European Patent No. 1 143 337 B, filed by Akamai Technologies, Inc. April 23, 2004 [43 pages].
- Reply of Dec. 21, 2004 of the Patent Proprietor to the Notice of Opposition of European Patent No. 1 143 337 B, Responsive to the Notice of Opposition of Akamai Technologies, Inc. dated April 23, 2004 [11 pages].
- Written Submissions in Opposition to European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 5, 2007 [5 pages].
- Communication from European Patent Office dated Feb. 8, 2007 regarding opposition to European Patent No. 1 143 337 B [11 pages].
- Owner's Response of March 20, 2008 to summons to attend oral proceedings dated January 24, 2008 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 3 to 7 [21 pages].
- Owner's Response of April 10, 2007 to summons to attend oral proceedings dated February 8, 2007 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 1 and 2 and attachment discussing claim features [26 pages].
- Claims of European Patent No. 1 143 337 B allowed/maintained in oral proceedings, April 24, 2008 [7 pages].

The Examiner is requested to contact the undersigned with any questions regarding the European patent or its opposition.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. No certification or fee is required.

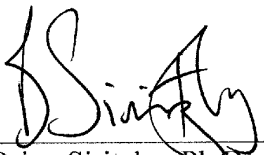
CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**

The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached

This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.

CUSTOMER NUMBER
42624

Respectfully submitted,

By: 

Brian Siritzky, Ph.D.
Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	1-1	Claims of European Patent No. 1 143 337 B allowed/maintained in oral proceedings, April 24, 2008 [7 pages].	
	1-2	Communication from European Patent Office dated Feb. 8, 2007 regarding opposition to European Patent No. 1 143 337 B [11 pages].	
	1-3	Notice of Opposition to a European Patent filed against European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 23, 2004 [5 pages].	
	1-4	Owner's Response of April 10, 2007 to summons to attend oral proceedings dated February 8, 2007 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 1 and 2 and attachment discussing claim features [26 pages].	
	1-5	Owner's Response of March 20, 2008 to summons to attend oral proceedings dated January 24, 2008 in Opposition Procedure of European Patent No. 1 143 337 B, including auxiliary requests 3 to 7 [21 pages].	
	1-6	Reply of Dec. 21, 2004 of the Patent Proprietor to the Notice of Opposition of European Patent No. 1 143 337 B, Responsive to the Notice of Opposition of Akamai Technologies, Inc. dated April 23, 2004 [11 pages].	
	1-7	Statement of Opposition to European Patent No. 1 143 337 B, filed by Akamai Technologies, Inc. April 23, 2004 [43 pages].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 2		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	2-1	Written Submissions in Opposition to European Patent No. 1 143 337 B, by Akamai Technologies, Inc., April 5, 2007 [5 pages].	
	2-2		
	2-3		
	2-4		
	2-5		
	2-6		
	2-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Electronic Acknowledgement Receipt

EFS ID:	3829591
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	42624
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	24-AUG-2008
Filing Date:	23-FEB-2005
Time Stamp:	15:17:20
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	NPL Documents	AK-written-submission-04-05-07.pdf	144333 78e53f4e8b89424eb7ae92d8ed66c8f6d9acef75	no	5

Warnings:	
Information:	

2	NPL Documents	Claims_maintained_in_opposition.pdf	149652 6355c6eff153e59b8f111fe593256d1c11957461	no	7
Warnings:					
Information:					
3	NPL Documents	Opposition_communication_from_EPO_02_08_07.pdf	472860 e9f801a95cbd14ee345eaec686b15e4653e37306	no	11
Warnings:					
Information:					
4	NPL Documents	Opposition_owner_10_04_07.pdf	911205 591a825df43ed57de438a49d2c93475a681b564b	no	26
Warnings:					
Information:					
5	NPL Documents	Opposition_papers.pdf	2154826 d5281a85ada5cacbbc2ded50099e27b63785dcff	no	21
Warnings:					
Information:					
6	NPL Documents	Opposition_Reply_of_Owner_12_21_04.pdf	356598 3f0d3e24ae5c44167cdcc41b05ec19cd17e4f726	no	11
Warnings:					
Information:					
7		Notice_of_Opposition_04_29_04.pdf	1809891 b94808a498aacf40971a7e531ae9703b3da063ee	yes	48

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	NPL Documents	1	5
	NPL Documents	6	48

Warnings:					
Information:					
8		20080824145532341.pdf	723833 57d9a4aeba8dcd1194811b73f020a37d7aa6f514	yes	5

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Information Disclosure Statement Letter	1	3

	Information Disclosure Statement (IDS) Filed (SB/08)	4	5
Warnings:			
Information:			
	Total Files Size (in bytes):	6723198	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.



JFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: August 21, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

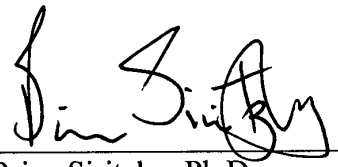
Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This IDS includes documents from litigation in related U.S. Patents Nos. 7,054,935 and 6,654,807. The PTO was previously advised of that litigation. The

Examiner is again requested to contact the undersigned with any questions about this litigation.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. No certification or fee is required.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 42624	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 • FAX: (703) 894-6430	



INFORMATION DISCLOSURE STATEMENT BY APPLICANT
FORM PTO-1449 (modified)

Sheet 1 of 7

Application No.	11/065,412
Filing Date	Feb. 23, 2005
First Named Inventor	Farber, David A.
Group Art Unit	2142
Examiner Name	BLAIR, Douglas B.
Attorney Docket No.	2711-0073
Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	1-1	US-5819092	10-06-1998	Ferguson
	1-2	US-6078943	06-20-2000	Yu
	1-3	US-6081835	06-27-2000	Antcliff
	1-4			
	1-5			
	1-6			
	1-7			
	1-8			
	1-9			
	1-10			
	1-11			
	1-12			
	1-13			
	1-14			
	1-15			
	1-16			
	1-17			
	1-18			
	1-19			
	1-20			
	1-21			
	1-22			
	1-23			
	1-24			
	1-25			
	1-26			

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	2-1	"Cisco DistributedDirector," Cisco Systems, Inc., White Paper, cited in Expert Report of Kevin Almeroth with date of 1996 (month unknown). [9 pages]	
	2-2	ALMEROTH, K., et al. "Scalable Delivery of Web Pages Using Cyclic Best-Effort (UDP) Multicast", IEEE INFOCOM, San Francisco, California, USA, June 1998.	
	2-3	ALMEROTH, K., Exhibit A to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [31 pages].	
	2-4	ALMEROTH, K., Exhibit B to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [131 pages].	
	2-5	ALMEROTH, K., Exhibit C to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [157 pages].	
	2-6	ALMEROTH, K., Exhibit D to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [128 pages].	
	2-7	ALMEROTH, K., Exhibit E Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [286 pages].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	3-1	ALMEROTH, K., Exhibit F to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [221 pages].	
	3-2	ALMEROTH, K., Exhibit G to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [369 pages].	
	3-3	ALMEROTH, K., Exhibit H to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [59 pages].	
	3-4	ALMEROTH, K., Exhibit I to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [95 pages].	
	3-5	ALMEROTH, K., Exhibit J to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [105 pages].	
	3-6	ALMEROTH, K., Exhibit K-1 to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008, [152 pages].	
	3-7	ALMEROTH, K., Exhibit K-2 to Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008 [25 pages].	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	4-1	ALMEROTH, K., Expert Report of Dr. Kevin C. Almeroth filed in Civil Action No. 2:07cv589 (RGD-FBS), U.S. Dist. Ct. E.D. Va., dated July 25, 2008. [64 pages and Exhibits A-J, K1, K2]	
	4-2	BAENTSCH, M., et al., "Introducing Application-Level Replication and Naming into Today's Web," Computer Networks and ISDN Systems, vol. 28, num. 7-11, pgs. 921-930, May 1996.	
	4-3	Berners-Lee, T. & Connolly, D., "Hypertext Markup Language - 2.0," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1866, Nov. 1995	
	4-4	BERNERS-LEE, T., et al., "Hypertext Transfer Protocol – HTTP/1.0," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-http-v10-spec-00.ps, March 1995.	
	4-5	BERNERS-LEE, T., et al., "Hypertext Transfer Protocol – HTTP/1.0," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1945, May 1996.	
	4-6	Berners-Lee, T.. "Resource Identifiers in WWW: A Unifying Syntax for the Expression of Names and Addresses of Objects on the Network as used in the World-Wide Web," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1630, June 1994.	
	4-7	Bestavros, A. Demand-based document dissemination to reduce traffic and balance load in distributed information systems. In Proc. IEEE Symp. on Parallel and Distributed Processing, San Antonio, TX, October 1995.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	5-1	CALLON, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1195, December 1990.	
	5-2	CHANDRA, R. et al., "BGP Communities Attribute," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1997, August 1996.	
	5-3	CHEN, E., et al., "An Application of the BGP Community Attribute in Multi-home Routing," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1998, August 1996.	
	5-4	COLAJANNI, M., et al., "Scheduling Algorithms for Distributed Web Servers," International Conf. on Distributed Computing Systems (ICDCS), Baltimore, Maryland, USA, May 1997.	
	5-5	Communication from European Patent Office dated May 30, 2008 in patent application EP 99906680 (12 pages).	
	5-6	ERIKSSON, H., "MBONE: The Multicast Backbone," Communications of the ACM, vol. 37, num. 8, Pg. 54-60, August 1994.	
	5-7	FIELDING, R., et al., "Hypertext Transfer Protocol – HTTP/1.1," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-http-v11-spec-00.txt, November 1995.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 6 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	6-1	FIELDING, R., et al., "Hypertext Transfer Protocol – HTTP/1.1," Internet Engineering Task Force(IETF) Request for Comments (RFC) 2068, January 1997.	
	6-2	MALKIN, G., "RIP Version 2 Carrying Additional Information," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1388, January 1993.	
	6-3	MOY, J., "OSPF Version 2," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1583, March 1994.	
	6-4	PARRIS C., et al, "A Dynamic Connection Management Scheme for Guaranteed Performance Services in Packet-Switching Integrated Services Networks," UC Berkeley Computer Science Division Tech. Report TR-93-005, 1993.	
	6-5	PARRIS C., et al, "The Dynamic Management of Guaranteed Performance Connections in Packet Switched Integrated Service Networks," UC Berkeley Computer Science Division and International Computer Science Institute Tech. Report CSD-94-859, 1994.	
	6-6	PARTRIDGE, C., et al., "Host Anycasting Service," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1546, November 1993.	
	6-7	POVEY, D. et al., "A distributed internet cache," in Proc. of the 20th Australasian Computer Science Conf., Feb. 1997.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 7 of 7		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	7-1	REKHTER Y., et al., "A Border Gateway Protocol 4 (BGP-4)," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1771, March 1995.	
	7-2	REKHTER Y., et al., "Application of the Border Gateway Protocol in the Internet," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1772, March 1995.	
	7-3	TOPOLCIC, C., "Experimental Internet Stream Protocol, Version 2 (ST-II)," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1190, October 1990.	
	7-4	TOUCH, J. et al., "The X-Bone,". Third Global Internet Mini-Conference at Globecom '98. Sydney, Australia, Nov. 1998, pp. 59-68 (pp. 44-52 of the mini-conference).	
	7-5	TRAINA, P., "BGP-4 Protocol Analysis," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1774, March 1995.	
	7-6	TRAINA, P., "Experience with the BGP-4 protocol," Internet Engineering Task Force (IETF) Request for Comments (RFC) 1773, March 1995.	
	7-7	WESSELS, D. et al., RFC 2187, "Application of Internet Cache Protocol (ICP), Version 2", Sep. 1997, Memo; Internet RFC/STD/FYI/BCP Archives; Network Working Group.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application Number:	11/065,412	Confirmation Number:	3296
First Named Inventor:	David A. Farber	Examiner:	BLAIR, DOUGLAS B
Filing Date:	02-23-2005	Attorney Docket Number:	2711-0073
Group Art Unit:	2142	Date:	August 5, 2008
Title:	**Shared Content Delivery Infrastructure**		

PRELIMINARY AMENDMENT

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Before initial examination, the Applicant respectfully requests the Examiner to include the following as though included in the initial filing.

Amendments to the Claims begin on page 2 of this paper.

Remarks begin on page 8 of this paper.

IN THE CLAIMS:

Please cancel claims 1-26 and add the following new claims:

27. (New) A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;

a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level domain name service (DNS) resolution;

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

28. (New) The distributed hosting framework as recited in claim 27, further including a redundant second level name server.

29. (New) A distributed hosting framework operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, the framework comprising:

at least one shared content server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers; and

at least one name server that provides domain name service (DNS) resolution,

wherein resources associated with the first content provider are associated with a first alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the first content provider made from said client machines; and

wherein the resources associated with the second content provider are associated with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources in response to requests for the resources associated with the second content provider made from said client machines.

30. (New) The distributed hosting framework as recited in claim 29, wherein the least one shared content server is addressable using the first alias name and the second alias name.

31. (New) The distributed hosting framework as recited in claim 29, further comprising:

a reflector constructed and adapted to associate the resources with the first alias name based, at least in part, on a relative cost of transmitting the resources to the client machines.

32. (New) The distributed hosting framework as recited in claim 29 further comprising:

a set of shared content servers that comprises the at least one content server; and

a reflector constructed and adapted to associate a particular resource with the first alias name after selecting the at least one shared content server from the set of shared content servers in response to a request for the particular resource issued by a particular client machine.

33. (New) The distributed hosting framework as recited in claim 29, wherein at least one of the resources is an embedded object.

34. (New) The distributed hosting framework as recited in claim 29, wherein the requests comprise host tag names that identify specific origin servers from which the resources originate.

35. (New) The distributed hosting framework as recited in claim 34, wherein the host tag names are included in HTTP headers.

36. (New) The distributed hosting framework as recited in claim 32, further comprising:

a load determination mechanism constructed and adapted to characterize load on the set of content servers for use by the reflector in selecting the at least one shared content server.

37. (New) The distributed hosting framework as recited in claim 36, wherein the reflector comprises the load mechanism.

38. (New) The distributed hosting framework as recited in claim 37, wherein the reflector is co-located with an origin server.

39. (New) The distributed hosting framework as recited in claim 36, further comprising:

a cost determination mechanism constructed and adapted to measure cost of delivering resources from each of the set of shared content servers to the particular client machine, the cost determination mechanism for use by the reflector in selecting the at least one shared content server.

40. (New) The distributed hosting framework as recited in claim 39, wherein the reflector comprises the cost determination mechanism.

41. (New) The distributed hosting framework as recited in claim 39, wherein the reflector is co-located with an origin server.

42. (New) The distributed hosting framework as recited in claim 39, wherein the cost comprises information relating a location associated with the particular client machine to each of the set of shared content servers.

43. (New) The distributed hosting framework as recited in claim 29, further including a redundant name server.

44. (New) The distributed hosting framework as recited in claim 29, wherein the at least one name server comprises:

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level DNS resolution.

45. (New) A method, in a distributed hosting framework operative in a computer network for delivering content to client machines and comprising a plurality of origin servers, each of said origin servers having resources associated therewith, at least one name server that provides domain name service (DNS) resolution, and at least one shared content server operable to replicate resources associated with the plurality of origin servers, the method comprising:

associating at least some resources associated with a first origin server of plurality of origin servers with a first alias name that causes the at least one name server to indentify the at least one shared content server to serve the resources associated with the first origin server in response to requests from client machines; and

associating at least some resources associated with a second origin server of said plurality of origin servers with a second alias name that causes the at least one name server to identify the at least one shared content server to serve the resources associated with the second origin server in response to requests from client machines.

46. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one shared content server from a set of shared content servers based at least in part on a network cost of delivering resources to the particular client machine.

47. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one content server from a set of shared content servers based at least in part on load characteristics of the set of shared content servers.

48. (New) The method of claim 45, wherein the associating step comprises:

in response to a request for a particular resource issued by a particular client machine, identifying the at least one content server from a set of shared content servers based at least in part on a location of the particular client machine relative to the at least one shared content server.

49. (New) The method as recited in claim 45, wherein the resources comprise an HTTP-based resource.

50. (New) The method as recited in claim 45, wherein the resources comprise a File Transfer Protocol (FTP)-based resource.

51. (New) The method of claim 45 wherein the at least one name server comprises:

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level DNS resolution.

REMARKS

By this Amendment, all pending claims have been canceled and new claims 27-51 have been added. Upon entry of this amendment, claims 27-51 will be pending in this application.

CLARIFICATION ON CLAIM NUMBERING

This application was filed with 24 claims. By a preliminary amendment with the filing of the application, claims 1-24 were canceled and new claim 25 was added. At that time claim 25 was the only pending claim. Then, on June 26, 2007, by another preliminary amendment, 26 new claims were added. These new claims were mistakenly numbered 2-26. Since the previously highest numbered claim was 25, these claims should have been numbered 26-47. However, the Office has entered these 25 claims as claims 2-26. Applicant has therefore numbered the new claims filed herewith as claims 27-52.

For clarification, the new claims filed herewith replace *all* previous claims. Should their numbering be incorrect, the Applicant respectfully requests the Examiner to renumber the claims or to indicate to the Applicant how the claims should be renumbered.

NOTE ON NEW CLAIMS 27 AND 28

New claims 27 and 28 correspond to claims 1 and 3 of **U.S. Patent No. 6,108,703**, Leighton *et al.,* filed May 19, 1999, issued August 22, 2000 ("the '703 Patent"). These claims of the '703 Patent copied herein were previously copied in application no. 09/612,598 on September 18, 2000, prior to one year from the date on which the '703 Patent was granted.

This application is a continuation of 09/612,598 which was a division of 09/021,506, filed February 10, 1998 (now patent no. 6,185,598 – "the '598 Patent"). The specification of this application is the same as that of the '598

Patent (except for the cross reference to related applications in the first paragraph of this application).

Applicant is <u>not</u> requesting an interference with claims 1 and 3 of the '703 Patent since those claims of the '703 Patent (corresponding to claims 27 and 28 hereof) have been invalidated in judicial proceedings, as being ***anticipated*** under 35 U.S.C. §102(e) by U.S. Patent No. 6,185,598, from which this application claims priority. *See Akamai Techs. v. Cable & Wireless Internet Services, Inc. et al.,* 344 F.3d 1186, 1197 (Fed. Cir. 2003), attached as Exhibit 1 hereto. As noted by the Federal Circuit in *Akamai Tech.*, "… we hold that claims 1 and 3 of the '703 patent are anticipated by the '598 patent and are invalid pursuant to 35 U.S.C. § 102." *Id.*

In view of their invalidation, claims 1 and 3 are no longer in the '703 Patent. Furthermore, in view of their invalidation as being *anticipated* by the '598 Patent, it follows that the instant application fully describes and enables those two claims.

CONCLUSION

Applicant respectfully submits that this application is in condition for allowance, and an early action allowing the claims is earnestly solicited. Should the Examiner believe that a telephone call will resolve any outstanding issues in this case, he is invited to telephone the undersigned at the number provided.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/ Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached. **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 42624	Respectfully submitted, By: /Brian Siritzky/37497 Brian Siritzky, Ph.D., Reg. No. 37497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 _ FAX: (703) 894-6430	

EXHIBIT 1

Akamai Techs. v. Cable & Wireless Internet Services, Inc. et al., 344 F.3d 1186 (Fed. Cir. 2003)

344 F.3d 1186
Akamai Technologies, Inc. and Massachusetts Institute of Technology, Plaintiffs-Appellees,
v.
Cable & Wireless Internet Services, Inc., Defendant-Appellant, and Kinetech, Inc., Defendant.
No. 03-1007.
United States Court of Appeals, Federal Circuit.
Decided: September 15, 2003.
Rehearing and Rehearing En Banc Denied November 13, 2003.
Page 1187
COPYRIGHT MATERIAL OMITTED
Page 1188

Mark T. Banner, Banner & Witcoff, Ltd., of Chicago, Illinois, argued for plaintiffs-appellees. With him on the brief were Pamela B. Krupka, of Washington, DC; and Aimee M. Boss, of Chicago, Illinois.

Arthur B. Wineburg, Pillsbury Winthrop LLP, of McLean, Virginia, argued for defendant-appellant. With him on the brief were Susan T. Brown, Daniel E. Yonan, of McLean, Virginia, and Brian J. Beatus, of Palo Alto, California. Of counsel were Brian Siritzky, Pillsbury Winthrop, LLP, of McLean, Virginia; and Eileen M. Herlihy and John T. Bennett, Palmer & Dodge LLP, of Boston, Massachusetts.

Before NEWMAN, GAJARSA, and DYK, Circuit Judges.

Opinion for the court by Circuit Judge GAJARSA. Opinion concurring in part and dissenting in part by Circuit Judge PAULINE NEWMAN.

GAJARSA, Circuit Judge.

Defendant-Appellant, Cable & Wireless Internet Services, Inc. ("C & W"), appeals the orders of the United States District Court for the District of Massachusetts: (1) denying C & W's motion for judgment as a matter of law challenging the jury verdict that claims 1, 3, 5, and 9 of U.S. Patent No. 6,108,703 ("the '703 patent") are valid and infringed; and (2) granting a permanent injunction based upon the jury verdict. Akamai Techs., Inc. v. Cable & Wireless Internet Serv., Inc., 2002 WL 1949231, No. 00-CV-11851 (D. Mass. May 24 and Aug. 21, 2002). Because claims 1 and 3, properly construed, are anticipated by U.S. Patent No. 6,185,598 ("the '598 patent") and are therefore invalid under 35 U.S.C. § 102, we reverse the district court's denial of C & W's motion for judgment as a matter of law with respect to claims 1 and 3 and vacate that portion of the permanent injunction. We affirm the district court's denial of C & W's motion for judgment as a matter of law with respect to the validity of claims 5 and 9 because those claims, properly construed, are not invalid. We also affirm the district court's denial of C & W's motion for judgment as a matter of law with respect to infringement of claim 9 because substantial evidence supports the jury's verdict of infringement. In view of our holding, we remand to the district court to review the scope of the permanent injunction.

I. BACKGROUND

The present appeal concerns technology for alleviating Internet congestion. To better understand the present dispute, a general overview of the relevant technology follows. Generally, people share information, i.e., "content," over the Internet through web pages. To look at web pages, a computer user accesses the Internet through a browser, e.g., Microsoft Internet Explorer® or Netscape Navigator®. These browsers display web pages stored on a network of servers commonly referred to as the Internet. To access the web pages, a computer user enters into the browser a web page address, or uniform resource locator ("URL"). The URL

fastcase

is typically a string of characters, e.g., www.fedcir.gov. This URL has a corresponding unique numerical address, e.g., 156.119.80.10, called an Internet Protocol ("IP") address. When a user enters a URL into the browser, a domain name service ("DNS") searches for the corresponding IP address to properly locate the web page to be displayed. The DNS is administered by a separate network of computers distributed throughout, and connected to, the Internet. These computers are commonly referred to as DNS servers. In short, a DNS server translates the URL into the proper IP address, thereby informing the user's computer where the host server for the web page www.fedcir.gov is located, a process commonly referred to as "resolving." The

Page 1189

user's computer then sends the web page request to the host server, or origin server. An origin server is a computer associated with the IP address that receives all web page requests and is responsible for responding to such requests. In the early stages of the Internet, the origin server was also the server that stored the actual web page in its entirety. Thus, in response to a request from a user, the origin server would provide the web page to the user's browser. Internet congestion problems quickly surfaced in this system when numerous requests for the same web page were received by the origin server at the same time.

This problem is exacerbated by the nature of web pages. A typical web page has a Hypertext Markup Language ("HTML") base document, or "container" document, with "embedded objects," such as graphics files, sound files, and text files. Embedded objects are separate digital computer files stored on servers that appear as part of the web page. These embedded objects must be requested from the origin server individually. Thus, each embedded object often has its own URL. To receive the entire web page, including the container document and the embedded objects, the user's web browser must request the web page and each embedded object. Thus, for example, if a particular web page has nine embedded objects, a web browser must make ten requests to receive the entire web page: one for the container document and nine for the embedded objects.

There have been numerous attempts to alleviate Internet congestion, including methods commonly referred to as "caching," "mirroring," and "redirection." "Caching" is a solution that stores web pages at various computers other than the origin server. When a request is made from a web browser, the cache computers intercept the request, facilitate retrieval of the web page from the origin server, and simultaneously save a copy of the web page on the cache computer. The next time a similar request is made, the cache computer, as opposed to the origin computer, can provide the web page to the user. "Mirroring" is another solution, similar to caching, except that the origin owner, or a third party, provides additional servers throughout the Internet that contain an exact copy of the entire web page located on the origin server. This allows a company, for example, to place servers in Europe to handle European Internet traffic.

"Redirection" is yet another solution in which the origin server, upon a request from a user, redirects the request to another server to handle the request. Redirection also often utilizes a process called "load balancing," or "server selection." Load balancing is often effected through a software package designed to locate the optimum origin servers and alternate servers for the quickest and most efficient delivery and display of the various container documents and embedded objects. Load balancing software locates the optimum server location based on criteria such as distance from the requesting location and congestion or traffic through the various servers.

Load balancing software was also known prior to the '703 patent. For example, Cisco Systems, Inc. marketed and sold a product by the name of "Distributed Director," which included server selection software that located the optimum server to provide requested information. The server selection software could be placed at either the DNS servers or the

content provider servers. The Distributed Director product was disclosed in a White Paper dated February 21, 1997 and in U.S. Patent No. 6,178,160 ("the '160 patent").

Page 1190

Both the White Paper and the '160 patent are prior art to the '703 patent. The Distributed Director product, however, utilized this software in conjunction with a mirroring system in which a particular provider's complete web page was simultaneously stored on a number of servers located in different locations throughout the Internet. Mirroring had many drawbacks, including the need to synchronize continuously the web page on the various servers throughout the network. This added extra expenses and contributed to congestion on the Internet.

Massachusetts Institute of Technology is the assignee of the '703 patent directed to a "global hosting system" and methods for decreasing congestion and delay in accessing web pages on the Internet. Akamai Technologies, Inc. is the exclusive licensee of the '703 patent.1 The '703 patent was filed on May 19, 1999, and issued on August 22, 2000. The '703 patent discloses and claims web page content delivery systems and methods utilizing separate sets of servers to provide various aspects of a single web page: a set of content provider servers (origin servers), and a set of alternate servers. The origin servers provide the container document, i.e., the standard aspects of a given web page that do not change frequently. The alternate servers provide the often changing embedded objects. The '703 patent also discloses use of a load balancing software package to locate the optimum origin servers and alternate servers for the quickest and most efficient delivery and display of the various container documents and embedded objects.

Independent claim 1, which is representative, reads:

1. A distributed hosting framework operative in a computer network in which users of client machines connect to a content provider server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname pretended to a domain name and path;

a set of content servers, distinct from the content provider server, for hosting at least some of the embedded objects of web pages that are normally hosted by the content provider server;

at least one first level name server that provides a first level domain name service (DNS) resolution; and at least one second level name server that provides a second level domain name service (DNS) resolution;

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

'703 patent, col. 17, ll. 17-37 (emphases added).

C & W is the owner, by assignment, of the '598 patent. The '598 patent is directed to similar systems and methods for increasing the accessibility of web pages on the Internet. The '598 patent was filed on February 10, 1998, and issued on February 6, 2001. Thus the '598 patent is prior art to the '703 patent pursuant to 35 U.S.C. § 102(e).2 C & W marketed and

Page 1191

sold products embodying the '598 patent under the name "Footprint." The relevant difference between the disclosure of the '598 patent and Akamai's preferred embodiment disclosed in the '703 patent is the location of the load balancing software. Akamai's preferred embodiment has the load balancing software installed at the DNS servers, while the '598 patent discloses installation of the load balancing software at the content provider, or origin, servers. The '598 patent does not disclose or fairly suggest that the load balancing software can be placed at the

DNS servers. It is now understood that placement of the software at the DNS servers allows for load balancing during the resolving process, resulting in a more efficient system for accessing the proper information from the two server networks. Indeed, C & W later created a new product, "Footprint 2.0," the systems subject to the permanent injunction, in which the load balancing software was installed at the DNS servers as opposed to the content provider servers. Footprint 2.0 replaced C & W's Footprint product.

On September 13, 2000, Akamai sued C & W seeking an injunction and damages for infringement of the '703 patent. Among other things, Akamai asserted that C & W's Footprint 2.0 content delivery network infringed apparatus claims 1, 3, 5, and 9 and method claims 17, 18, and 22 of the '703 patent.

C & W answered the Complaint alleging that Footprint 2.0 did not infringe the claims of the '703 patent and that the claims of the '703 patent were invalid under 35 U.S.C. §§ 102 and 103(a). In particular, C & W asserted that the '598 patent anticipated the asserted claims and that the asserted claims were obvious in light of the '598 patent in combination with Cisco's Distributed Director product.

The district court conducted a Markman hearing and entered its order construing the disputed terms of the '703 patent (as well as two other patents that are not at issue in this appeal). The district court stated that "[a]t the time of the Markman hearing, the parties ... did not appear to have a common understanding as to which additional claims were still in dispute. The parties' written and oral presentations offer little assistance in this regard. I therefore limit my ruling to the terms above." Notably, the parties did not appear to dispute the construction of any terms in the "wherein" clause of claim 1.

After a 19-day jury trial, the jury determined that C & W infringed apparatus claims 1, 3, 5, and 9 and method claims 17, 18, and 22. The jury upheld the validity of apparatus claims 1, 3, 5, and 9. The jury invalidated claims 17, 18, and 22 under 35 U.S.C. §§ 102 and 103(a) as either anticipated by the '598 patent or obvious in light of the '598 patent in view of Cisco's Distributed Director product. C & W filed its motion for judgment as a matter of law in February 2002 asserting that claims 1, 3, 5, and 9 were invalid and/or not infringed. The district court denied that motion and permanently enjoined C & W from "making, using, selling, offering for sale, or importing into the United States the patented inventions claimed in claims 1, 3, 5, and 9 of the '703 patent," in an Order that "extended to Footprint 2.0 service, as configured and described at trial."

C & W appealed the denial of its motion for judgment as a matter of law and the district court's entry of a permanent injunction on the bases that claim 9 was not infringed and that claims 1, 3, 5 and 9 were invalid for anticipation and/or obviousness. C & W did not appeal the infringement of claims 1, 3, and 5. We have jurisdiction

Page 1192

over this appeal pursuant to 28 U.S.C. §§ 1292(a), (c)(1), and 1295(a)(1).

II. DISCUSSION

We review the denial of a motion for judgment as a matter of law following a jury verdict by reapplying the district court's standard of review. Catalina Lighting, Inc. v. Lamps Plus, Inc., 295 F.3d 1277, 1284 (Fed.Cir.2002); Stryker Corp. v. Davol, Inc., 234 F.3d 1252, 1257 (Fed.Cir.2000); Tec Air, Inc. v. Denso Mfg., 192 F.3d 1353, 1357 (Fed.Cir.1999). Thus, we review claim construction, an issue of law, de novo. SIBIA Neurosciences, Inc. v. Cadus Pharm. Corp., 225 F.3d 1349, 1354 (Fed.Cir.2000).

With regard to factual findings, we must presume that the jury resolved all factual disputes in favor of the prevailing party, and we must leave those findings undisturbed as long as they are supported by substantial evidence. Id. A factual finding is supported by substantial evidence if a reasonable jury could have found

in favor of the prevailing party in light of the evidence presented at trial. Tec Air, 192 F.3d at 1358. "Substantial evidence is more than a mere scintilla. It means such relevant evidence as a reasonable mind might accept as adequate to support a conclusion." Consol. Edison Co. v. NLRB, 305 U.S. 197, 229, 59 S.Ct. 206, 83 L.Ed. 126 (1938). Thus, substantial evidence review involves an examination of the record as a whole, taking into consideration evidence that both justifies and detracts from the decision of the fact-finder. In re Gartside, 203 F.3d 1305, 1312 (Fed. Cir.2000); Nat'l Presto Indus., Inc. v. W. Bend Co., 76 F.3d 1185, 1192 (Fed.Cir.1996) (holding that a jury verdict must be sustained if it is supported by substantial evidence based on a review of the entirety of the record). In reviewing the record, we must draw all reasonable inferences in favor of the prevailing party, and not make credibility determinations or substitute our view of the conflicting evidence for that of the jury. SIBIA, 225 F.3d at 1355 (citing Connell v. Sears, Roebuck & Co., 722 F.2d 1542, 1546 (Fed.Cir.1983)). If, however, after reviewing all of the evidence in a light most favorable to the prevailing party, this court is convinced that a reasonable jury could not have found in that party's favor, we must reverse the denial of a motion for judgment as a matter of law. Id.

A. Anticipation

C & W appeals the jury finding that claims 1 and 3 were valid as not anticipated by the disclosure of the '598 patent. The first step in any invalidity analysis is claim construction, an issue of law this court reviews de novo. Cybor Corp. v. FAS Techs., Inc., 138 F.3d 1448, 1456 (Fed.Cir.1998) (en banc). The second step, determining whether a prior art reference discloses each and every limitation of the claim expressly or inherently, Scripps Clinic & Research Found. v. Genentech, Inc., 927 F.2d 1565, 1576-77 (Fed. Cir.1991), is a factual question reviewed for substantial evidence. Eaton Corp. v. Rockwell Int'l Corp., 323 F.3d 1332, 1343 (Fed.Cir.2003). This factual question is contingent upon the proper claim construction. Id. at 1344. A claim limitation is inherent in the prior art if it is necessarily present in the prior art, not merely probably or possibly present. Rosco v. Mirror Lite, 304 F.3d 1373, 1380 (Fed.Cir.2002). "[T]he dispositive question regarding anticipation is whether one skilled in the art would reasonably understand or infer from the prior art reference's teaching that every claim [limitation] was disclosed in that single reference." Dayco Prods., Inc. v. Total Containment, Inc., 329 F.3d 1358,

Page 1193

1368 (Fed.Cir.2003) (internal quotation marks and alterations omitted).

Through trial and on appeal, the parties have narrowed the disputed issues of validity to a single point of contention — the placement of the load balancing software at either the DNS servers or the origin server. Therefore, our initial focus in the anticipation analysis is on the construction of claims 1 and 3, in particular whether claims 1 and 3 require the presence of load balancing software at the DNS servers. The issue before us is thus a relatively self-contained one. On the one hand, if claims 1 and 3 require load balancing at the DNS servers, the claims are not anticipated. On the other hand, if the claims do not require this limitation, they are anticipated by the '598 patent. The only disputed limitation of claims 1 and 3 reads:

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

'703 patent, col. 17, ll. 31-37 (emphases added).

Claim 3 is dependent upon independent claim 1 and includes the following additional limitation.

3. The hosting framework as described in claim 1 further including a redundant second level name server.

Id., col. 17, ll. 40-41.

Akamai contends that the '598 patent differs from claims 1 and 3 of the '703 patent in the placement of the load balancing software. Indeed, in its brief on appeal, Akamai stated:

The significant difference between the prior art '598 patent and the '703 patent claims on appeal was acknowledged and admitted by everyone throughout the trial.... In particular, C & W counsel told the jury the difference involves the fact that selection of the best computer server to deliver the embedded objects of the web page is done in the '598 prior art patent by "software ... located at the origin server" whereas selection of the best computer server to deliver the content is done in the '703 patent "by software located at the DNS...."

C & W argues that the location of the load balancing software is not a limitation in claims 1 and 3, and in the alternative, that while the '598 patent does not explicitly disclose the placement of load balancing software at the DNS servers, it is nevertheless inherent in the Internet and the '598 patent.

We agree that claims 1 and 3 do not include a load balancing limitation. While the written description unquestionably contemplates the preferred location of the load balancing software, claims 1 and 3 do not expressly require its presence. To support its reading of independent claim 1, Akamai points only to the term "identifying" in the "wherein" clause of claim 1 which states:

wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the content provider server and the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers.

This language, however, requires only that the embedded object is "identified by the modified embedded object URL" and is "served from a given one of the content

Page 1194

servers as identified by the first and second level name servers." The plain meaning of the claim language does not require any load balancing mechanism. Instead, it simply requires the embedded object to be served from "the content servers as identified by the first level and second level name servers." Load balancing, if required at all, could be at either the DNS servers or the content provider server. The ordinary meaning of the term "identifying" in claims 1 and 3 covers standard DNS resolution, without any sort of load balancing. Absent evidence that a "patentee unequivocably imparted a novel meaning to [the] term[ ] or expressly relinquished claim scope during prosecution," we give the limitation its full ordinary and customary meaning. Omega Eng'g, Inc. v. Raytek Corp., 334 F.3d 1314, 1323 (Fed. Cir.2003); Teleflex, Inc. v. Ficosa N. Am. Corp., 299 F.3d 1313, 1325-26 (Fed.Cir.2002); CCS Fitness, Inc. v. Brunswick Corp., 288 F.3d 1359, 1366-67 (Fed.Cir.2002); Renishaw PLC v. Marposs Societa' per Azioni, 158 F.3d 1243, 1249 (Fed.Cir.1998) ("Absent a special and particular definition created by the patent applicant, terms in a claim are to be given their ordinary and accustomed meaning.").

The only question that remains is whether the written description or the prosecution history unequivocally shows that the inventors imparted a novel meaning to the term "identifying" to include load balancing. Omega Eng'g, 334 F.3d at 1323; Teleflex, 299 F.3d at 1325-26. The written description does not specifically define the term "identifying." Rather, the discussion with respect to load balancing focuses on the DNS servers as performing "special function[s]," e.g., load balancing functions, without any reference to the term "identifying." See, e.g., '703 patent, col. 9, ll. 31-48. Similarly, the parties have pointed to nothing in the prosecution history with respect to the term

"identifying." Akamai's only evidence that supports its special definition of the term "identifying" is the testimony of one of the inventors, Mr. Farber, of the '598 patent. Mr. Farber stated that: "the DNS in our system are a little different because we did the step of identifying ... which repeater should be used by the browser as part of the, using [sic] the HTTP method instead of the DNS method."3 This extrinsic evidence is not the unequivocal evidence, Omega Eng'g, 334 F.3d at 1323, indicating the term "identifying" should take anything other than its ordinary and accustomed meaning. While this possibly suggests that the inventors believed the "identifying" step included a load balancing function, "what the patentee subjectively intended his claims to mean is largely irrelevant to the claim's objective meaning and scope." Solomon v. Kimberly-Clark Corp., 216 F.3d 1372, 1379 (Fed. Cir.2000). It is also not testimony that clearly supports the proposition that the term "identifying" has a special meaning to one of ordinary skill in the art.

Thus claim 1, as properly construed, does not include the limitation of the placement of the load balancing mechanism. The parties agree that the '598 patent discloses all the remaining limitations of claim 1. Because claim 1 does not require exact placement, it is therefore invalid as anticipated by the '598 patent.

Claim 3 similarly does not require placement of the load balancing software at the DNS servers. Claim 3 only includes the additional limitation that the hosting framework as described in claim 1 further includes "a redundant second level name

Page 1195

server." On appeal, Akamai's primary argument echoes that of claim 1, namely that the load balancing software is located at the DNS servers. Akamai's only separate argument with respect to claim 3 is that "[b]ecause [the] '598 patent did not even mention hierarchical DNS (i.e., more than one level), clearly the jury was entitled to reject the notion that [the] '598 [patent] also anticipated claim 3." This additional argument, however, fails to address C & W's contention that hierarchical DNS is inherent in any Internet system. Indeed, C & W proffered documentary evidence and testimony at trial that redundant domain name servers are inherent in any Internet-based application. See Dayco, 329 F.3d at 1369. Akamai points to no evidence whatsoever that contradicts the evidence presented to the jury at trial. Accordingly, we hold that any inference in favor of Akamai relating to the redundant second level server in claim 3 is unsupported by substantial evidence. The addition of a redundant second level DNS server does not save the validity of dependent claim 3. Claim 3 is therefore also invalid under 35 U.S.C. § 102 as anticipated by the '598 patent.4

B. Obviousness

C & W next claims that it is entitled to judgment as a matter of law that claims 5 and 9 are obvious in light of the '598 patent in combination with Cisco's Distributed Director product and are therefore invalid.5 Claims 5 and 9, which are dependent on independent claim 1, include the following additional limitations.

5. The hosting framework as described in claim 1 wherein the second level name server includes a load balancing mechanism that balances loads across a subset of the set of servers.

9. The hosting framework as described in claim 1 wherein the first level name server includes a network map for use in directing a request for the embedded object generated by a client.

We review the ultimate determination of obviousness de novo. Modine Mfg. Co. v. Allen Group, Inc., 917 F.2d 538, 541 (Fed.Cir.1990). This ultimate determination, however, requires underlying factual findings, which this court reviews to determine whether they are supported by substantial evidence and, if they are, whether those findings support the legal conclusions which necessarily were drawn by the jury in forming its verdict. Id. A claimed invention is

unpatentable due to obviousness if the differences between it and the prior art "are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art." 35 U.S.C. § 103(a) (2000). While the ultimate conclusion of obviousness is for the court to decide as a matter of law, several factual inquiries underlie this determination. Graham v. John Deere Co., 383 U.S. 1, 17-18, 86 S.Ct. 684, 15 L.Ed.2d 545 (1966). These inquiries include the scope and content of the prior art, the level of ordinary skill in the field of the

Page 1196

invention, the differences between the claimed invention and the prior art, and any objective evidence of nonobviousness such as long-felt need and commercial success. Id. When a rejection depends on a combination of prior art references, there must be some teaching, suggestion, or motivation to combine the references. In re Geiger, 815 F.2d 686, 688 (Fed.Cir.1987). Although the suggestion to combine references may flow from the nature of the problem, Pro-Mold & Tool Co. v. Great Lakes Plastics, Inc., 75 F.3d 1568, 1573 (Fed.Cir.1996), the suggestion more often comes from the teachings of the pertinent references, In re Sernaker, 702 F.2d 989, 994 (Fed.Cir.1983), or from the ordinary knowledge of those skilled in the art that certain references are of special importance in a particular field, Pro-Mold, 75 F.3d at 1573 (citing Ashland Oil, Inc. v. Delta Resins & Refractories, Inc., 776 F.2d 281, 297 n. 24 (Fed.Cir.1985)). Therefore, "[w]hen determining the patentability of a claimed invention which combines two known elements, `the question is whether there is something in the prior art as a whole to suggest the desirability, and thus the obviousness, of making the combination.'" In re Beattie, 974 F.2d 1309, 1311-12 (Fed.Cir.1992) (quoting Lindemann Maschinenfabrik GMBH v. Am. Hoist & Derrick Co., 730 F.2d 1452, 1462 (Fed.Cir.1984)).

To prevail, C & W must therefore show that no reasonable jury could have found claims 5 and 9 nonobvious in light of the evidence presented. Tec Air, 192 F.3d at 1358. Here, C & W has not met this burden. Namely, we are unable to discern any suggestion or motivation to combine the references as C & W suggests in the record before us. Indeed, when pressed on this issue at oral argument, C & W could point only to Cisco's Distributed Director product, which at best disclosed that load balancing software could be placed at either the DNS servers or the origin servers for a "mirroring" system. In its brief, C & W pointed to the fact that the '598 patent, the '703 patent, and Cisco's Distributed Director product all address the same problem: Internet congestion.

Our review of the evidence presented to the jury does not persuade us that no reasonable jury could have found claims 5 and 9 nonobvious over the prior art, and we decline to disturb the factual findings of the jury. Drawing all factual inferences in favor of Akamai, we affirm the jury's verdict with respect to the validity of claims 5 and 9 because there was no suggestion or motivation to combine the references.

Additionally, the record contains substantial evidence relating to secondary considerations supporting the jury's verdict. In particular, the record shows that C & W expended significant effort to determine how Akamai's products worked. Once it was determined that Akamai placed the server selection software at the DNS servers, C & W redesigned its Footprint product, abandoning the embodiments in the '598 patent. The new Footprint 2.0 design incorporated Akamai's placement of the load balancing mechanism at the DNS server. This evidence of copying is relevant to an obviousness determination. See Advanced Display Sys., Inc. v. Kent State Univ., 212 F.3d 1272, 1285 (Fed.Cir.2000); Vandenberg v. Dairy Equip. Co., 740 F.2d 1560, 1567 (Fed.Cir.1984) ("The copying of an invention may constitute evidence that the invention was not an obvious one.... This would be particularly true where the copyist had itself attempted for a substantial length of time to design a similar device, and had failed."). C & W's redesign process was documented in the

record in internal emails from C & W engineers discussing

Page 1197

Akamai's approach, identifying weaknesses in C & W's approach, and ultimately deciding to switch to the Akamai system.

In sum, C & W has not shown that the jury's conclusion that claims 5 and 9 of the '703 patent are nonobvious under 35 U.S.C. § 103(a) is unsupported by substantial evidence. Accordingly, we hold that claims 5 and 9 of the '703 patent are not obvious in light of the '598 patent in view of Cisco's Distributed Director product, and therefore affirm the district court's ruling denying C & W's motion for judgment as a matter of law and affirm the grant of a permanent injunction against C & W with respect to claims 5 and 9.

C. Infringement of Claim 9

Having addressed validity of the various claims of the '703 patent, we next turn to the jury's finding of infringement with respect to dependent claim 9. We review the record as a whole to determine whether substantial evidence exists to support the jury's finding that claim 9 is infringed. C & W objects to the jury verdict based on the fact that Dr. Bustavros's expert report was limited to infringement under the doctrine of equivalents. Akamai counters with cites to the testimony of Dr. Bustavros, who testified that C & W's Footprint 2.0 system included a network map routine and therefore infringed claim 9, without any reference to a limitation under the doctrine of equivalents. This dispute appears to be one of interpretation and weight of a witness's testimony — one for the jury to resolve. This court generally does not tread on the jury's role in making these determinations. Thus, we affirm the jury's finding of infringement of claim 9 as supported by substantial evidence, namely Dr. Bustavros's testimony at trial.

III. CONCLUSION

For the foregoing reasons, we hold that claims 1 and 3 of the '703 patent are anticipated by the '598 patent and are invalid pursuant to 35 U.S.C. § 102. We therefore reverse the district court's denial of C & W's motion for judgment as a matter of law with respect to the validity of claims 1 and 3, and instruct the district court to modify the permanent injunction accordingly. We also affirm the district court's denial of C & W's motion for judgment as a matter of law with respect to claims 5 and 9 and therefore affirm the district court's issuance of the permanent injunction with respect to those claims. We remand to the district court to review and redetermine the scope of the permanent injunction.

AFFIRMED-IN-PART, REVERSED-IN-PART AND REMANDED.

IV. COSTS

No costs.

Notes:

1. For purposes of this opinion both plaintiffs are collectively referred to as "Akamai."

2. Akamai does not dispute that the '598 patent is "prior art" to the '703 patent for purposes of validity under 35 U.S.C. § 102 or § 103.

3. The "HTTP method" refers to placement of the load balancing software at the origin servers responsible for providing the HTTP container page.

4. Judge Newman in her dissent specifically points out that the issue of anticipation is a question of fact. Claim construction, however, is a question of law. Before the factual question of anticipation may be addressed, a court must first properly construe the claims before it. Here, claims 1 and 3 were not properly construed by the district court. Therefore, a necessary first step in this court's anticipation inquiry was to properly construe the claims at issue.

5. C & W also challenges the validity of claims 1 and 3 under the same obviousness theory. Because we hold that claims 1 and 3 are anticipated by the '598 patent, we need not reach this argument.

PAULINE NEWMAN, Circuit Judge, concurring in part, dissenting in part.

I concur in much of the court's decision. However, the court provides no sufficient basis for overturning the findings of the jury and overturning the affirmance by the district court with respect to the question of anticipation of claims 1 and 3.

The issue of anticipation is a question of fact, and the jury verdict that there is not anticipation must be sustained if there is substantial evidence in its support. It is not disputed that the prior art (the defendant's '598 patent) does not disclose and does not embody the DNS lookup of the Akamai '703 patent. There was extensive evidence, presented by both sides, as to the content of the prior art; there was substantial evidence that the prior art's origin server "reflector" is a different structure and method, and that the subject

Page 1198

matter of claims 1 and 3 does not read on the prior art. If there were any question concerning claims 1 and 3, the question would be one of obviousness; not anticipation. The jury verdict that these claims were not anticipated was supported by substantial evidence, with a lengthy and thorough and fully presented trial, and a verdict that could have been reached by a reasonable jury. The criteria of reversal are not met.

Akamai correctly points out that the defendant makes no effort to discuss the support, or lack thereof, for the jury verdict. My colleagues on this panel commit the same error, for the majority opinion says not a word about the evidence at trial, but simply decides the question for itself. Reversal of the judgment rendered on a jury verdict is appropriate only if there is no legally sufficient evidentiary basis for the verdict. See Fed.R.Civ.P. 50(a)(1); Intercity Maintenance Co. v. Local 254, Service Employees International Union, 241 F.3d 82, 86 (1st Cir.2001); National Presto Industries, Inc. v. West Bend Co., 76 F.3d 1185, 1192 (Fed.Cir.1996).

The sufficiency of the evidence must be reviewed in the light most favorable to the party that received the verdict, with all reasonable inferences drawn in favor of the verdict. Sheils Title Co. v. Commonwealth Land Title Insurance Co., 184 F.3d 10, 19 (1st. Cir.1999); SIBIA Neurosciences, Inc. v. Cadus Pharmaceutical Corp., 225 F.3d 1349, 1355 (Fed.Cir.2000). It was not disputed, indeed it was admitted by C & W, that the prior art did not show either the claim 1 limitation that the content server is "distinct from the content provider server" (clause 3 of claim 1) or that "the embedded object identified by the modified embedded object URL is served from a given one of the content servers as identified by the first level and second level name servers" (sixth clause of claim 1). Even C & W's technical expert Dr. Dewar conceded that "the one thing that is not taught explicitly by the '598 patent is the use of DNS." Claim 1's clauses 4 and 5 require "first level" and "second level" "domain name service (DNS) resolution." Although Dr. Dewar went on to express the opinion that these changes would have been obvious, he did not testify that they were present, even inherently, in the '598 reference. The prior art was not shown to suggest or use or contemplate the DNS as used by the MIT inventors.

Claim 1 requires identification of the content server of the distributed hosting framework during the DNS lookup. Witnesses for both sides agreed that this differs from ordinary use in the context of the internet, and that persons experienced in this field would readily so understand. It is seriously incorrect for this court to reconstruct the invention and then to invalidate the claims on its own findings, ignoring the evidence at the trial.6

There was substantial evidence by which the jury could have found that all of the limitations of claims 1 and 3 are not present in the prior art. The acknowledged differences from the prior art render untenable the panel majority's restatement of the issues, as well as their resolution of the factual question of anticipation in order to invalidate the claims.

I respectfully dissent.

Notes:

6. The claims were construed by agreement before trial, and the only issues relate to validity and infringement. It is inappropriate to recharacterize the factual question of anticipation as one of claim construction, instead of deciding the appeal on the basis on which it was tried, and on the appropriate standard for review of jury verdicts.

Electronic Acknowledgement Receipt

EFS ID:	3734407
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	42624
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	05-AUG-2008
Filing Date:	23-FEB-2005
Time Stamp:	19:39:37
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes) /Message Digest	Multi Part /.zip	Pages (if appl.)
1	Preliminary Amendment	2711_0073_Prelim_Amdt.pdf	395048 103cda532a601da2b7400b3c1fb58e2e6c7e0aae	no	22
Warnings:					
Information:					

Total Files Size (in bytes):	395048

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (07-06)
Approved for use through 1/31/2007. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
11/065,412	02/23/2005	

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)	SMALL ENTITY ☐		OR	OTHER THAN SMALL ENTITY	
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =		OR	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =			X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)	SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
	08/05/2008	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	* 25	Minus	** 25	=	X $ =		OR	X $ =	
	Independent (37 CFR 1.16(h))	* 3	Minus	***	=	X $ =		OR	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1)		(Column 2)	(Column 3)					
		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)		RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =		OR	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =		OR	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))									
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".

*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

Legal Instrument Examiner:
/FREDERICK E. BRISCOE/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: June 6, 2008

INFORMATION DISCLOSURE STATEMENT

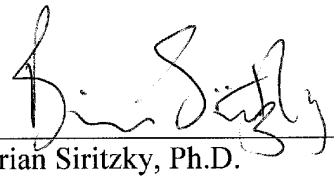
Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed within three (3) months of the U.S. filing date OR before the mailing date of a first Office Action on the merits. No certification or fee is required.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**	
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**	
CUSTOMER NUMBER 42624	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 • FAX: (703) 894-6430	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 1 of 1		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	1-1	Korkea-aho, M. (1995). Scalability in Distributed Multimedia Systems, Technical report TKO-B128, Helsinki University of Technology.	
	1-2		
	1-3		
	1-4		
	1-5		
	1-6		
	1-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Electronic Acknowledgement Receipt

EFS ID:	3416795
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	42624
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	06-JUN-2008
Filing Date:	23-FEB-2005
Time Stamp:	15:07:46
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes) /Message Digest	Multi Part /.zip	Pages (if appl.)
1		20080606145151972.pdf	355924 87430d8ea3fb3874350f01cd04647e64d0a3cec0	yes	3

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Information Disclosure Statement Letter	1	2
	Information Disclosure Statement (IDS) Filed	3	3

Warnings:					
Information:					
2	NPL Documents	korkea-aho95scalability.pdf	1373753 6efedf610476e5dc0f76a1e73cf75d839ac4ce1f	no	96
Warnings:					
Information:					
		Total Files Size (in bytes):	1729677		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IFW



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: May 12, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached. It is respectfully requested that the information be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed within three (3) months of the U.S. filing date OR before the mailing date of a first Office Action on the merits. No certification or fee is required.

Copies of some of the references were cited by or submitted to the Office in parent application No. 09/612,598 filed July 7, 2000, which is relied upon for an earlier filing date under 35 USC 120. Thus form PTO 1449 is attached without copies of these references. 37 CFR 1.98(d).

RELATED LITIGATION

The Examiner is reminded that this patent application is related to U.S. Patent Nos. 7,054,935 and 6,654,807, both of which have been asserted in a litigation for patent infringement in the Eastern District of Virginia, in a case captioned: *Level 3 Communications, LLC v. Limelight Networks, Inc.* (2:07-cv-00589-RGD-FBS). The case was filed on 12/17/2007. Documents relating to that litigation were cited in a IDS filed January 25, 2008.

The Examiner is again respectfully requested to contact the undersigned with any questions about the related litigation.

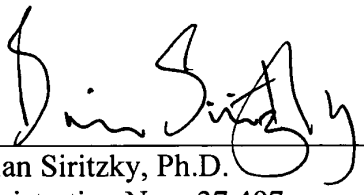
CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**

The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached

This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.

CUSTOMER NUMBER
42624

Respectfully submitted,

By: 

Brian Siritzky, Ph.D.
Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

OIPE
MAY 12 2008
PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 5	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	1-1	US-5708780	01/13/98	Levergood et al.
	1-2	US-5826031	10/1998	Nielsen
	1-3	US-5838906	11/1998	Doyle et al.
	1-4	US-5931904	1999-08-03	Banga
	1-5	US-5944780	1999-08-31	Chase
	1-6	US-5973696	10-26-99	Agranat, et al.
	1-7	US-5996025	11/30/99	Day
	1-8	US-6006264	12/1999	Colby et al.
	1-9	US-6021426	2000-02-01	Douglis
	1-10	US-6041307	March 21, 2000	Ahuja, et al.
	1-11	US-6098078	2000-08-01	Gehani
	1-12	US-6108703	08/2000	Leighton et al.
	1-13	US-6115752	09/2000	Chauhan
	1-14	US-6125394	2000-09-26	Rabinovich
	1-15	US-6144996	11/2000	Starnes et al.
	1-16	US-6154738	11/2000	Call
	1-17	US-6154777	11/28/2000	Ebrahim
	1-18	US-6185598	02/2001	Farber
	1-19	US-6189039	02/13/01	Harvey
	1-20	US-6195680	02/27/01	Goldszmidt
	1-21	US-6226618	05/01/01	Downs
	1-22	US-6263313	07/17/01	Milsted
	1-23	US-6272566	08/07/01	Craft
	1-24	US-6282574	08/28/2001	Voit
	1-25	US-6311214	10/2001	Rhoads
	1-26	US-6330602	Dec-11-01	Law et al.

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 5	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	2-1	US-6332195	12/2001	Green et al.
	2-2	US-6398245	06/04/02	Gruse
	2-3	US-6418421	07/09/02	Hurtado
	2-4	US-6442549	08/2002	Schneider
	2-5	US-6460082	10/01/02	Lumelsky
	2-6	US-6463454	10/08/02	Lumelsky
	2-7	US-6463508	10/08/02	Wolf
	2-8	US-6470389	Oct-22-02	Chung et al
	2-9	US-6587837	07/01/03	Spagna
	2-10	US-6654807	25-Nov-03	Farber et al.
	2-11	US-6763377	07/13/04	Belknap et al
	2-12	US-6859791	02/22/05	Spagna
	2-13	US-6963910	11/08/05	Belknap
	2-14	US-6973485	12/2005	Ebata et al.
	2-15	US-7103564	09/05/06	Ehnebuske
	2-16	US-7110984	09/19/06	Spagna
	2-17	US-7117259	10/03/06	Rohwer
	2-18	US-7188085	03/06/07	Pelletier
	2-19	US-7206748	04/17/07	Gruse
	2-20			
	2-21			
	2-22			
	2-23			
	2-24			
	2-25			
	2-26			

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication Date	Name of Patentee or Applicant of Cited Document	Notes
	3-1	CA 2202572	10/14/1998	Nortel Networks Ltd.	
	3-2	EP-0824236 A2	02/18/1998	Baird	
	3-3	EP-0865180 A2	09/16/1998	Lucent	
	3-4				
	3-5				
	3-6				
	3-7				
	3-8				
	3-9				
	3-10				
	3-11				
	3-12				
	3-13				
	3-14				
	3-15				
	3-16				
	3-17				
	3-18				
	3-19				
	3-20				
	3-21				
	3-22				
	3-23				
	3-24				
	3-25				
	3-26				

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	4-1	Awerbuch, B. et al., Distributed Paging for General Networks. In Proc. of the 7th ACM-SIAM Symposium on Discrete Algorithms, pp. 574-583, January 1996.	
	4-2	Baentsch, M, et al. "Enhancing the Web's Infrastructure: From Caching to Replication." IEEE Internet Computing, 1(2):18--27, Mar. 1997.	
	4-3	CARTER, J. Lawrence et al., "Universal Classes of Hash Functions", Journal of Computer and System Sciences, Vol. 18, No. 2, (April 1979), 143-154	
	4-4	Cisco Systems, Inc., "Cache Director System", http://www.cisco.com/warp/public/751/cache/cdsov.htm, (1997), 1-5	
	4-5	Cisco Systems, Inc., "Cisco Cache Engine", http://www.cisco.com/warp/public/751/cache/index.html	
	4-6	Gadde, S., et al., "Reduce, reuse, recycle: An approach to building large internet caches," in Workshop on Hot Topics in Operating Systems (HotOS), April 1997.	
	4-7	Naor, Moni, et al., "The Load, Capacity and Avaliability of Quorum Systems", In Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, (November 1994), 214-225	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 5		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	Notes
	5-1	Pultar, Giray; "USENIX, Automatically Selecting a Close Mirror Based on Network Topology," Proceedings of the Twelfth Systems Administraion Conference; Boston, Massachusetts, Dec. 1998, p. 159-166.	
	5-2	Ravi, R., "Rapid Rumor Ramification: Approximating the Minimum Broadcast Time", In Proceedings of teh 35th IEEE Symposim on Foundation of Computer Science, (November 1994), 202-213	
	5-3	Schmidt, J., et al., "Chernoff-Hoeffding Bounds for Applications with Limited Independence", in Proc. 4th ACM-SIAM Symp. on Discrete Algorithms, (1993), 331-340.	
	5-4	WEGMAN, Mark, et al., "New Hash Functions and Their Use in Authentication and Set Equality", Journal of Computer and System Sciences, Vol. 22, (June 1981), 265 -279	
	5-5		
	5-6		
	5-7		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant. Notes: If identified, the following is provided: EA = English Abtract, T = Translation, PF = Patent Family.

Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2202572 C 2004/02/10

(11)(21) 2 202 572

(12) BREVET CANADIEN
CANADIAN PATENT

(13) C

(22) Date de dépôt/Filing Date: 1997/04/14

(41) Mise à la disp. pub./Open to Public Insp.: 1998/10/14

(45) Date de délivrance/Issue Date: 2004/02/10

(51) Cl.Int.^{6}/Int.Cl.^{6} G06F 15/16, G06F 13/38, H04L 29/02

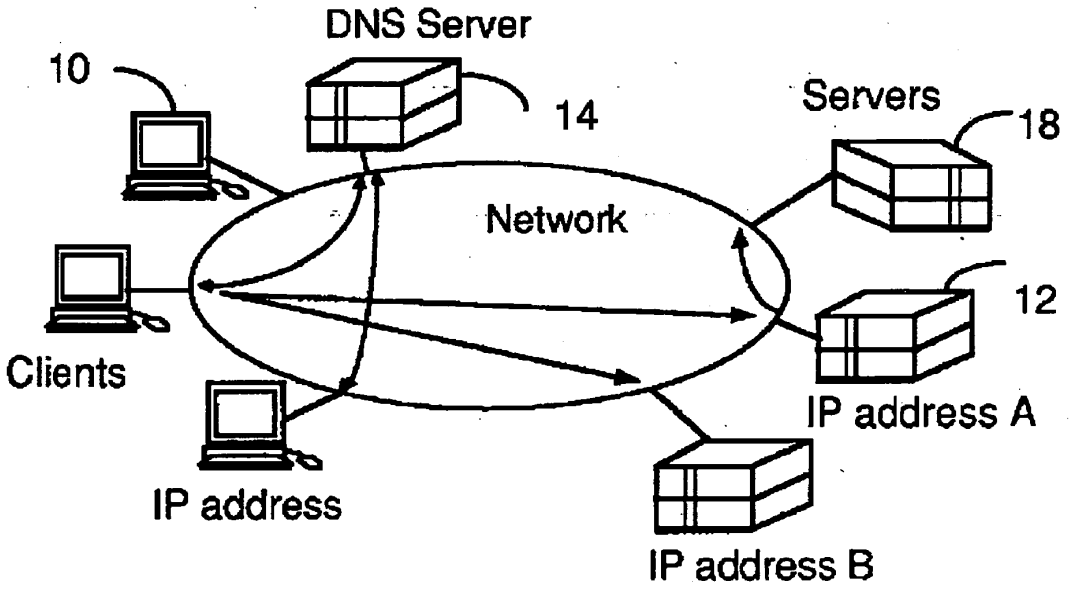
(72) Inventeurs/Inventors:
LAW, KA LUN EDDIE, CA;
NANDY, BISWAJIT, CA;
CHAPMAN, ALAN STANLEY JOHN, CA

(73) Propriétaire/Owner:
NORTEL NETWORKS LIMITED, CA

(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : SERVEUR WEB A VARIABILITE DIMENSIONNELLE ET METHODE DE GESTION EFFICACE DE SERVEURS MULTIPLES

(54) Title: A SCALEABLE WEB SERVER AND METHOD OF EFFICIENTLY MANAGING MULTIPLE SERVERS



(57) Abrégé/Abstract:
A network includes one or more clients, a plurality of servers and an intermediate device called a "depot" Information resources are replicated among the servers The client accesses a server through the depot The depot transparently identifies a new session and selects one server for the session The depot may be provided between an Internet Client and an ATM network server. An edge node may be provided between the depot and the ATM server A proxy may be provided between the depot and a WWW server.

Abstract

A network includes one or more clients, a plurality of servers and an intermediate device called a "depot". Information resources are replicated among the servers. The client accesses a server through the depot. The depot transparently identifies a new session and selects one server for the session. The depot may be provided between an Internet Client and an ATM network server. An edge node may be provided between the depot and the ATM server. A proxy may be provided between the depot and a WWW server.

A SCALEABLE WEB SERVER AND METHOD OF EFFICIENTLY MANAGING MULTIPLE SERVERS

Field of the Invention

The invention relates generally to distributed server systems which are used widely in the field of telecommunications for information sharing. In particular, it is directed to a distributed server system which is scaleable and realizes high server performances. The distributed server system of the invention includes an intermediary node between a client and a pool of servers.

Background of the Invention

The last few years have observed a phenomenal growth in web (short for World-Wide Web or WWW, or Internet) usage. This growth has demonstrated the value of wide-area information-sharing but, at the same time, caused a significant research interest in improving the performance of web systems. Recent studies show that the web consumes more Internet bandwidth than any other application.

At a macro level, a web system consists of three components: (i) client, (ii) communication protocol, and (iii) server. Efforts are being made at each component to enhance the performance of the overall web systems.

At the client end, supports are provided to improve response time by the following features: memory cache, disk cache, allowing multiple simultaneous sessions and introducing a proxy server for another level of caching.

The communication protocol between the client and the server is HTTP (Hypertext Transfer Protocol) which always assumes the existence of a reliable path layer underneath the client and server. TCP/IP (Transmission Control Protocol/Internet Protocol) provides reliable data transmission using window flow control techniques. HTTP therefore runs on top of the TCP/IP layer. Asynchronous Transfer Mode (ATM) is another transmission technique to handle broadband multimedia traffic. It continues to grow steadily in the communication world. High-speed ATM switches are available in the commercial market. The co-existence of the Internet and large-scale ATM networks

is expected in the near future. ATM can provide wide-area virtual circuits, thus facilitating geographical distribution of web servers.

The HTTP has also undergone changes for performance improvement. It has been reported that multiple TCP sessions per HTTP transaction is a major cause of performance bottleneck. The introduction of a "keep alive" header allows sessions to be kept open and used for multiple HTTP request/response activities.

The web servers have also undergone improvements: first-generation servers handled 20 transactions/sec. based on one process per transaction. The major overhead was due to a large process fork time for a new transaction. This was avoided by pre-forking multiple processes and using a dispatcher to distribute transactions among them to achieve server performance of 100 transactions/sec. The "keep alive" HTTP feature, along with the multi-threaded architecture of only one process, allows the server to handle more than 250 transactions/sec. However, the current trend indicates that the popular sites will incur a significantly higher number of server transactions per second in the near future. This requires more powerful web servers which may be developed by improving different components of a web server (e.g., CPU speed, disk performance, file system performance, performance of TCP/IP, server software architecture etc.). Alternatively, multiple servers can be used to handle high rate of server transactions.

The multiple server approach has two immediate advantages: if a server fails, the session can be handled from other servers; also the total cost of multiple servers can be less than the cost of one server with the equivalent performance. It is therefore foreseen that multiple server systems will be in great demand to accommodate an ever-increasing number of user transactions.

Different architectures for multiple server systems are currently in use and are described briefly here. The use of a Domain Name System (DNS) server to distribute traffic among multiple servers was investigated at NCSA at the University of Illinois. Figure 1 shows the DNS system. When a client 10 wishes to communicate with a server 12, at first it contacts the DNS 14, from which it obtains the IP address of the desired server. The client then uses this IP address to communicate with the server. All clients perform the same process unless they

already have the IP addresses of servers with which they want to connect. When there are a plurality of servers which hold identical information, the DNS rotates in a round robin manner through a pool of these identical servers which are alternatively mapped to the alias of the hostname of one server. This approach has provided some success in distributing the server load, however, it could not balance the load among servers. Another problem with this approach is that, once the IP address resolution is cached in the local memory, the client may never contact DNS.

Another system uses the HTTP level redirection capability to move a transaction among multiple servers. Figure 1 also shows this mechanism. When a server 12 finds that it is impossible to handle any extra traffic, it can redirect a transaction to another preselected server 18 and hence distribute the load. HTTP redirection is a common technique used for WWW load distribution. The implementation maybe simple and straightforward, but the redirection requires a round trip delay between the client and server before the transaction is redirected to a different server. Moreover, if the first server is already very busy, the response delay will be even greater.

Figure 2 shows another known system which switches the load based on the client IP address. Each client 20 goes to an intermediate device 22 which examines the originating IP address and decides where to forward the traffic among multiple servers 24. IP address hashing is one of the possible mechanisms to determine the server to which the traffic will be directed. This technique, however, lacks the dynamic control of user accesses. Moreover, the IP address spaces are partitioned into five different classes. Care should therefore be taken in designing good hashing function.

HTTP is a stateless protocol. A web server obtains everything it needs to know about a request from the HTTP request itself. After the request is serviced, the server can forget the transaction. Thus, each request in HTTP is disjointed. If all the servers are identical (or see the same file system using a distributed file system), the server from which the request is served is of little relevance to the client. The choice of a physical server itself is immaterial to the transactions. An HTTP transaction is an aggregation of one or more TCP sessions. Based on this principle, different TCP sessions can be allocated to different

servers without the knowledge of whether or not all the TCP sessions belong to the same HTTP transaction. The present invention realizes this TCP-based switching by the use of an intermediary entity called a depot to perform these functions of session allocation. Thus, the TCP-based server switching allows a nice granularity for load balancing among multiple servers. It is also envisaged that this concept of forwarding different sessions to different servers can be applied to similar multi-server architectures of telecommunications networks.

Objects of the Invention

It is an object of the invention to provide a server system which has a high performance.

It is another object of the invention to provide a communications architecture which allows the construction of a scaleable and high performance server system.

It is a further object of the invention to provide a method of and an apparatus for efficiently managing the resources of a multiple server system.

It is still another object of the invention to provide a method of and an apparatus for efficiently managing the resources of a multiple server system by the use of an intermediate entity.

It is yet a further object of the invention to provide a method of and an apparatus for efficiently managing the resources of a multiple server system by the use of more than one instance of the intermediate entity for improved robustness.

Summary of the invention

According to an aspect of the present invention, there is provided a method of efficiently utilizing the plurality of servers in a client-server environment of a telecommunications network, where information resources, which are accessed by HTTP protocol, are replicated among the servers and a client has access to any one or more of the servers for desired information resources. The method includes the steps of: performing an HTTP transaction by HTTP protocol between the client and the plurality of servers, the HTTP transaction comprising one or more TCP

connection sessions which are conducted by way of a depot; storing, at the depot, information about the TCP connection sessions and parameters in a primary map and a secondary map, the information in the primary map being deleted at the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL); transparently forwarding TCP connection sessions of the HTTP transaction at the depot in both directions to and from the client and the plurality of servers; monitoring parameters indicative of either or both of load and operability of the servers; detecting transition states of the TCP connection sessions of the HTTP transaction at the depot; and switching at the depot the plurality of TCP connection sessions of the HTTP transaction among the plurality of servers based on the detected transition states and the monitored parameters of the servers so that during the HTTP transaction, transfer of the information resources is performed by HTTP protocol between the client and one or more servers among the plurality of servers.

According to an aspect of the present invention, there is provided a method of efficiently utilizing the plurality of servers in a client-server environment of a telecommunications network, where information resources, which are accessed by HTTP protocol, are replicated among the servers and a plurality of clients have access to any one or more of the servers for desired information resources through a plurality of HTTP transactions. Each HTTP transaction uses multiple TCP connection sessions. The method includes the steps of: performing each HTTP transaction on HTTP protocol between any one of the plurality of clients and the plurality of servers, each HTTP transaction comprising one or more TCP connection sessions which are conducted on TCP/IP through a depot. The method further includes the following steps performed at the depot: storing information about the TCP connection sessions and parameters in a primary map and a secondary map, the information in the primary map being deleted at

the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL); inspecting all TCP/IP packets to detect a state transition of the TCP connection sessions; transparently forwarding TCP/IP packets of existing TCP connection sessions to a correct server; detecting TCP/IP packets of a new TCP connection session; monitoring parameters indicative of either one or both of load and operability of the servers; selecting a server among the plurality of servers based on their monitored parameters and the state transition; and transparently forwarding all the TCP/IP packets of the new TCP connection session to the selected server so that during each of the existing and new TCP connection sessions, transfer of the information resources of the HTTP transactions is performed between any one of the clients and one or more servers by one or more new and existing TCP connection sessions.

According to an aspect of the present invention, a depot for transferring one or more TCP connection sessions between a client and among a plurality of servers in a client-server environment of a telecommunications network, where information resources are replicated among the servers and a plurality of clients have access to any one or more of the servers for desired information resources through an HTTP transaction using at least two layers of protocols, one stateless protocol layer upon another stateful protocol layer. The HTTP transaction runs on the stateless protocol layer of HTTP protocol and is made up of one or more TCP connection sessions, which run on TCP/IP layer. The depot includes: a primary and a secondary map for storing information about the TCP connection sessions and parameters, the information at the primary map being deleted at the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL); a packet analysis block for analysing and transparently forwarding TCP/IP packets of one or more TCP connection sessions of the HTTP transaction between the

client and one of the servers; a session management block for switching one or more TCP connection sessions of the HTTP transaction among the plurality of servers according to a preset criteria including the TCP/IP states of each TCP connection session; and a protocol state tracking block for monitoring the TCP/IP states of the TCP connection sessions of the HTTP transaction.

Brief Description of the Drawings

Figure 1 shows multiple server systems of known architecture;

Figure 2 shows another multiple server system of known architecture;

Figure 3 shows schematically a client-server server system according to one embodiment of the invention;

Figure 4 is a functional block diagram of an intermediary device according to one embodiment of the invention;

Figure 5 is a schematic illustration of the invention as applied in the ATM environment according to one embodiment;

Figure 6 shows protocol stacks of the architecture shown in Figure 5;

Figure 7 is a functional block diagram of the embodiment shown in Figure 5;

Figure 8 is a schematic illustration of a yet another embodiment of the invention; and

Figure 9 shows schematically a yet further embodiment of the invention as applied in the proxy server environment.

Detailed Description of Preferred Embodiments of the Invention

Figure 3 shows schematically the invention according to one embodiment. A client 30 is communicating with a server 32 through a telecommunications network which can be any transport network such as telephone networks, data networks, ATM networks, LANs, or a combination thereof. WWW or Internet runs on top of the transport network. Typically, a client sends the server 32 a request for information stored therein and the server replies with the desired information. The desired information, however, may require more than one session (e.g., TCP session). In the multiple server environment, servers 34 and 36 contain identically stored information as that stored in server 32. According to the one embodiment of invention, an intermediate entity 38 sits transparently between the client and the servers. The intermediate entity 38 is called a "depot" throughout this specification and is the heart of the invention. When the depot receives a request for information, it distributes the TCP sessions among multiple servers based on the server load balancing criteria. In actual implementation, the depot can be realized by the use of appropriate software, hardware or a combination of the both. It should also be noted that the concept of session distribution among the servers can be implemented in other multiserver environments which use protocol stacks of similar configurations.

The depot is a forwarding mechanism and forwards any particular packet to the assigned server using a map. All clients access the server system using the depot's IP address. The distribution of TCP sessions among multiple servers by the depot remains transparent to the client. A TCP session consists of multiple TCP packets. All TCP packets of a given TCP session are served by the same server. Since the forwarding of all TCP packets of a given TCP session must go to the same server, a map between the IP address and port number of the

client and the identity of the server is maintained. The entry of this map is to be maintained in the depot as long as required by the TCP protocol. This entry is preserved in the depot as long as there is a possibility of the arrival of a TCP packet from a client or a server. The depot does not generate any TCP packets nor does it discard any unless they are anomalous enough that there is no way to assign them to a server.

A depot has the following functions:

1. inspect all packets in both directions at IP and TCP levels;
2. choose a server based on load balancing criteria for a new TCP session;
3. forward TCP packets for existing sessions to the already chosen server;
4. forward TCP packets from servers to clients;
5. clean up the mapping entry when TCP sessions end; and
6. watch for and handle anomalous TCP packets.

All these functions are performed in a depot and Figure 4 shows its functional block diagram. Referring to Figure 4, the depot includes TCP packet forwarding block 40 for TCP packet analysis and forwarding. This block inspects all packets in both directions at IP and TCP levels, forwards TCP packets of existing sessions to the already chosen server, and forwards TCP packets from servers to the clients. For a new session, this block identifies the TCP session setup request and forwards information to session management block 42. The session management block 42 chooses a server among a pool of available servers. This block is preloaded with the next server to be allocated. This pre-load is done by running a load balancing algorithm such as shown at 44 in the figure on the basis of knowledge of the server statistics and perhaps network states, or simply by round-robin. Server probing 48 can be performed periodically to obtain the server statistics. The TCP packets carrying data are forwarded by the TCP packet forwarding block 40 by identifying the entry of a map between the server and the client. At TCP state tracking block 46, TCP states of both the clients and the servers are being monitored to facilitate the management of the TCP session, e.g., opening/closing and control functions.

The depot maintains the map 50 of all the active TCP sessions in a table called the primary table. Each session between a client and a server is identified by the combination of the client's IP address and port number.

A typical primary table entry has the following: client IP address and port number, server identity, TCP states and related parameters (ack number, seq number etc.). The depot maintains another list of sessions in a table called the secondary table. The entries of the primary table are deleted at the close of sessions under normal or anomalous conditions. The secondary table entries are maintained for a 2 MSL (maximum segment lifetime) period. The duration of a 2 MSL differs for different implementations but, in this embodiment, a period of 2 minutes is chosen as an example. The 2 MSL is important since a TCP packet may arrive after a session is closed, due to variable delay at the network.

The reason for maintaining two separate tables at the depot is to reduce the search space for finding a match on arrival of a TCP packet. The number of entries in the secondary table is very large since the entries are maintained for a 2 MSL period. The entries in the primary table are for active TCP sessions only. Most of the incoming TCP packets should find a match in the primary table.

A TCP session is stateful within each session between a client and a server. Therefore, the depot should handle all arriving TCP/IP packets as transparently as possible; otherwise, the introduction of an intermediary such as the depot may disrupt the normal TCP logic between the two end systems. The crucial function of the state tracking routine in the depot involves the creation and deletion of the table entries. Several types of TCP/IP packets can initiate the creation and deletion of table entries. In the TCP state transition diagram, there are eleven states. For the client-server model, however, the number of possible states is reduced. This reduces the state tracking requirement in the depot. However, out-of-sequence situations can occur in the Internet, and different permutations of state transitions can arise which slightly increase the complexity of the state tracking routine. Moreover, in wide-area networks, packet loss is possible. Therefore, any packet received by the depot does not guarantee the reception of the packet by the other end. As a result, according to the invention,

packets in both directions are analysed in order to get a complete view of the state of the session and the state tracking function of the depot simply guesses the current TCP states of the client and server by investigating the "flag" information of all arriving TCP packets. The states guessed at the depot may be different from the actual states at the client and server, if a TCP packet is dropped or corrupted in the network between the depot and the client or server.

The following is a detailed description of the handling of various TCP packets according to one embodiment of the invention.

SYN: A SYN packet header (client IP address and port) is matched with the entries in the primary table and then with the secondary table. If no match is found (i.e., arrival of a new session), a new entry is created in the primary table and based on the load balancing criterion, a server is allocated for the session. If a match is found (i.e., duplicate SYN), the packet is forwarded to an already allocated server.

FIN, PSH, URG: The match is found from the primary or secondary table and forwarded to the appropriate server. If no match is found, the packet is dropped.

ACK: All ACK packets are forwarded to the server if an entry is found in the table. If the ACK packet causes the state transition to the TIME_WAIT state, the entry is moved from the primary table to the secondary table for a 2 MSL time-out.

RST: All reset packets are validated by checking if the sequence number is in the window. If the state is SYN-SENT, the RST is considered valid if the ACK field acknowledges the SYN. If the RST is valid, the entry from the primary table is moved to the secondary table.

The exception conditions are handled in the following manner.

If an entry in the primary table is inactive for a long time (e.g., more than 20 mins.), the entry is moved to the secondary table. This is necessary since a client or server may crash without proper termination of a TCP session, causing an entry to the primary table to remain forever.

If an entry in the tables (primary and secondary) is not found on arrival of a packet other than SYN, the packet is dropped. This is necessary since the depot does not know the destination for the packet.

The depot does not cause any interruption of the ongoing TCP sessions between the clients and the servers. This is because the depot forwards every packet, irrespective of the states guessed at the depot, as long as a table entry is found. For example, a RST may be lost in the network between the depot and client but the entry is still maintained in the secondary table, which enables the forwarding of any retransmitted RST. In one embodiment, the forwarding from the depot to the servers is achieved by directing packets to physical ports of the device containing the depot function where each physical port connects to one server.

According to a further embodiment, the depot system of the invention is implemented in the ATM environment. In this embodiment, an ATM network is used as a transport mechanism to provide distributed WWW services. The network model is shown in Figure 5. On the ATM network, multiple, information homogeneous, but potentially geographically distributed web servers 50 are provided. An Internet web client 52 can access information via the routes shown on the diagram. A depot 54 of the invention provides the interface between an Internet client and an ATM network server. An incoming packet from the client will be segmented and encapsulated into AAL5 cells at the depot without the requirement of modifying either the IP or TCP headers. The cells (packets) are carried sequentially in a virtual channel.

Like the earlier embodiment, the depot ideally sits transparently between the clients accessing from the Internet and the servers on the ATM networks. The depot employs a TCP-based switching mechanism. It examines all TCP/IP packets and sends them to the ATM network from the Internet and vice versa. Figure 6 shows the protocol stack structure of the system model. On the depot, there are TCP and IP stacks but the TCP session is only observed and not terminated.

Therefore, the depot of the embodiment performs the same functions as those described earlier. In addition, however, it must perform the following functions:

- transform IP packets to ATM cells and vice versa;
- map the QoS parameters from IP to ATM where applicable; and
- participate as necessary in the ATM network management.

Similar to Figure 4, Figure 7 shows a functional block diagram of the depot of this embodiment. The depot identifies the flow (client IP and port) from the TCP/IP packets and maps the corresponding VPI/VCI in its table entries. At the IP/ATM adaptation block, the depot performs segmentation and encapsulation of packets into AAL5 cells without the requirement of modifying either the IP or TCP headers. Likewise, it also perform the reverse functions of ATM cells from the ATM network. For a new session, the packet analyzer identifies the TCP session setup request and forwards the information to the session management block. The next server to be allocated is selected, as described earlier. Packets are kept contiguous and are all sent over the same VC.

If the session is already allocated and recognized by the packet analyzer, it will read the correct VPI/VCI from its tables and forward the packet onward. Packets in the reverse direction are also analyzed in order to have a complete view of the state of the session.

In another embodiment, the forwarding is achieved by encapsulating the IP frames in a frame transport, such as Ethernet or Frame relay, using the addressing of that frame transport to identify individual servers.

In another embodiment, the forwarding is achieved by modifying the header information of the TCP/IP packet and redirecting it over an IP network.

In Figure 8, another embodiment is shown. According to this embodiment, servers may or may not be ATM-based hosts and therefore an edge node 80 is required as an access point to each server, which provides traffic management and any necessary IP/ATM conversion. Figure 8 shows all the servers and Internet connections accessing the ATM network via edge nodes. The edge nodes communicate with each other over independent virtual circuits. Therefore, depot 82 is also an edge node managing the Internet connection that has the additional function of routing incoming packets to the correct server.

According to a yet further embodiment of the invention, the depot system of server management is applied to the proxy server architecture. Proxy servers are used to reduce the network load and latency by migrating the work load close to the clients.

A proxy server is an application gateway which operates at the HyperText Transfer Protocol (HTTP) layer. The basic function of a proxy server is almost identical to a HTTP server in transferring client requested documents. Furthermore, it is able to interpret and modify a HTTP request before forwarding it to the end WWW server. This is because the proxy server has the caching function. If there is a cache hit, it delivers the found documents to the client locally, thus reducing the network load and transmission latency.

Figure 9 shows schematically the concept of another embodiment of the invention as applied to the proxy server management system. In Figure 9, a depot proxy system is located between an intranet and the Internet. The depot distributes sessions among a pool of proxy servers based on load balancing or other criteria. The functions of the depot of the proxy system is identical to those described earlier. Therefore, for a new session, the packet analyzer identifies the TCP session setup request and forwards the information to the session management block. If the session is already allocated and recognized by the packet analyzer, it will read in the correct proxy identity from its tables and forward the packet onward. Packets in the reverse direction are also analyzed in order to have a complete view of the state of the session.

This proxy architecture can achieve the following similar goals:

1. scaleable proxy server arrangements;
2. high availability of information service; and
3. dynamic load balancing of traffic loads to different proxy servers.

To achieve scaleability, more proxy servers can be attached to the depot. This is because the depot handles only simple functions. When a depot operates to its limitations, it is also possible to add a new depot with a new IP address under the same alias name. With the newly added depot, another cluster of proxy servers can be created on the network.

Experimental Results

A Pentium PC running a NetBSD operating system is used for the software implementation of the depot. The client's HTTP requests are generated using benchmark software from Zeus Corporation. The clients and servers are connected to the depot using 10 Mbps Ethernet.

The same test file is retrieved 1000 times for each test. The number of simultaneous HTTP requests are varied to study the depot behaviour. Traffic is forwarded by the depot to two identical NCSA/1.5.1 Web servers in a round robin fashion. Server throughput (total bytes transferred per second) is measured for each test case. The total number of bytes of data and the HTTP headers divided by the time taken to transfer indicates the server throughput. The number of HTTP requests served by the server per second is also measured. Both measures include a variable network delay. To minimise the variable network delay, the experiment is performed at a time when the network is very lightly loaded. The variation of the server throughput and HTTP requests served per second with an increased rate of client requests are good indicators of server performance.

Table 1

	file size: 100 bytes		file size: 1 Kbytes		file size: 10 Kbytes	
Concurrent sessions	single server	invention	single server	invention	single server	invention
1	18.22	17.46	66.58	56.67	52.13	51.98
5	29.21	44.62	99.26	80.75	151.85	168.18
10	21.1	55.13	94.83	99.81	226.80	244.53
15	10.18	39.5	65.75	101.73	223.1	277.06
20	11.17	29.86	42.64	111.93	251.24	289.91
25	6.46	30.00	33.86	117.53	233.73	297.89
100	6.55	12.23	X	37.59	X	282.81

Table 2

	file size: 100 bytes		file size: 1 Kbytes		file size: 10 Kbytes	
Concurrent sessions	single server	invention	single server	invention	single server	invention
1	64.61	61.92	55.16	46.95	5.00	4.99
5	103.48	158.23	82.19	66.88	14.55	16.13
10	74.59	194.70	78.31	82.55	21.68	23.44
15	35.73	138.95	54.03	84.21	21.35	26.53
20	39.03	104.62	34.87	92.29	23.97	27.67
25	22.48	105.22	27.55	97.23	22.40	28.49
100	22.41	40.79	X	28.51	X	25.69

Table 1 shows the server throughput in Kbytes/sec. The experiment is performed with three different file sizes: 100 bytes, 1 Kbyte and 10 Kbytes. The number of simultaneous client HTTP requests are varied from 1 to 100. The total number of client HTIP requests for each test is 1000. The experiment is performed with a single server system and a depot system with two servers. Table 2 shows the number of served HTTP requests per second under the same experimental setup.

For one client HTTP request (any file size) at a time, the depot system is always slower than the single server. This is due to the store and forward delay introduced at the depot. For multiple simultaneous HTTP requests, however, the HTTP requests are served in parallel by two servers. In general, the depot system shows improved throughput and served HTTP requests per second. Another reason for performance improvement is the reduced workload at each server due to the distribution of HTTP requests. For 100 bytes and 1 Kbyte files, the single server performance degrades quickly with the increased number of simultaneous sessions. Performance with the depot system is consistently better than the single server system. For the 10 Kbytes file, the single server performance remains flat. However, the performance with the depot system is better than the single server system.

For 100 simultaneous HITP requests, the single server system could complete the test for only a 100 bytes file size. The depot system could complete the tests for all three file sizes.

It is found that the throughput and HTTP requests served per second with the depot system are more than double those of the single server system with a large number of simultaneous sessions. For example, the number of HTTP requests served per second for the test case of 1 Kbyte file size with 20 simultaneous sessions is 92.29. This is an improvement of 2.65 over the single server system. This linear improvement is due to the halved number of simultaneous sessions (i.e., 10) on each server. The number of HTTP requests served per second by the single server system with 10 simultaneous HTTP requests is 78.31. Thus, maximum served requests of (2*78.31) 156.62 is theoretically attainable.

Thus it is shown that the server throughput and serviced requests per second can be improved using multiple servers according to the invention.

In the multiple server system, if one server goes down, the depot transfers all subsequent sessions to the other servers in the same cluster. There are several methods to detect server failure. A depot may assume the failure of a server if it does not receive an ACK from the server for several seconds after it sends a message to the server. After that, the depot starts polling the failed server on a regular basis to check if the server program has restarted. Another method that the depot uses is the passive reception of server statistics as an indication of the operability of a server program. If the depot does not receive any information for a period of time, then the depot will put this server into a dead-list and new incoming TCP requests will not be routed to this server. When server statistics are received from this server again, the depot will then put it back into an alive-list. However, there is difficulty in identifying the current state within each TCP session such as the number of bytes sent, the current window size and sequence number etc. Thus, the lost TCP sessions from a faulty server are simply scrapped.

The depot reliability is also very important since all the traffic is concentrated at the depot. Standby sparing is a simple method of providing fault tolerance in the system. There can be hot or warm standby. A sparing standby depot machine is always in alert mode to assume normal depot operations upon detecting a failure in the operating depot. For hot standby, the replica actively monitors all input streams from both the client and server. Both operating machine and replica have virtually identical TCP state information about all sessions. Detection of depot failure initiates changeover of control.

For warm standby, the operating machine periodically sends the state information to the replica to update the standby's database. In this situation, the two databases shall not be the same most of the time, but the goal is to reduce the number of affected TCP sessions when failure occurs. Then the standby replica will automatically be activated on fault detection.

What is claimed is:

1. In a client-server environment of a telecommunications network, where information resources, which are accessed by HTTP protocol, are replicated among a plurality of servers and a client has access to any one or more of the servers for desired information resources, a method of efficiently utilizing the plurality of servers comprising steps of:

performing an HTTP transaction by HTTP protocol between the client and the plurality of servers, the HTTP transaction comprising one or more TCP connection sessions which are conducted by way of a depot;

storing, at the depot, information about the TCP connection sessions and parameters in a primary map and a secondary map, the information in the primary map being deleted at the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL);

transparently forwarding TCP connection sessions of the HTTP transaction at the depot in both directions to and from the client and the plurality of servers;

monitoring parameters indicative of either or both of load and operability of the servers;

detecting transition states of the TCP connection sessions of the HTTP transaction at the depot; and

switching at the depot the plurality of TCP connection sessions of the HTTP transaction among the plurality of servers based on the detected transition states and the monitored parameters of the servers so that during the HTTP transaction, transfer of the information resources is performed by HTTP protocol between the client and one or more servers among the plurality of servers.

2. The method according to claim 1 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

3. The method according to claim 1 wherein the depot is connected to one or more servers through an ATM network, comprising a further step of:

performing the TCP connection sessions in ATM cells through the ATM network.

4. The method according to claim 3 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

5. The method according to claim 1 wherein the depot is connected to one or more servers through a local area network or a frame relay network, comprising further steps of:

performing the TCP connection sessions in their appropriate format through the local area network or the frame relay network.

6. The method according to claim 5 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

7. The method according to claim 1, wherein the depot is connected to one or more servers by way of a plurality of proxy servers, comprising further steps of:

switching at the depot the TCP connection sessions among the proxy servers which in turn send the HTTP transaction to the servers; and

caching information resources contained in the HTTP transaction at the proxy servers.

8. The method according to claim 7, comprising further steps of:

monitoring the status of all the proxy servers; and

for each TCP connection session, selecting one proxy server among all the proxy servers based on their monitored status.

9. The method according to claim 8 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

10. In a client-server environment of a telecommunications network, where information resources, which are accessed by HTTP protocol, are replicated among a plurality of servers and a plurality of clients have access to any one or more of the servers for desired information resources through a plurality of HTTP transactions, each HTTP transaction using multiple TCP connection sessions, a method of efficiently utilizing the plurality of servers during the HTTP transactions, comprising steps of:

performing each HTTP transaction on HTTP protocol between any one of the plurality of clients and the plurality of servers, each HTTP transaction comprising one or more TCP connection sessions which are conducted on TCP/IP through a depot; and

at the depot, performing the following steps:

storing information about the TCP connection sessions and parameters in a primary map and a secondary map, the

information in the primary map being deleted at the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL);

inspecting all TCP/IP packets to detect a state transition of the TCP connection sessions;

transparently forwarding TCP/IP packets of existing TCP connection sessions to a correct server;

detecting TCP/IP packets of a new TCP connection session;

monitoring parameters indicative of either one or both of load and operability of the servers;

selecting a server among the plurality of servers based on their monitored parameters and the state transition; and

transparently forwarding all the TCP/IP packets of the new TCP connection session to the selected server so that during each of the existing and new TCP connection sessions, transfer of the information resources of the HTTP transactions is performed between any one of the clients and one or more servers by one or more new and existing TCP connection sessions.

11. The method according to claim 10 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of the operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

12. The method according to claim 10 wherein the depot is connected to one or more servers through an ATM network, comprising a further step of:

performing the TCP connection sessions in ATM cells through the ATM network.

13. The method according to claim 12 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of the operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

14. The method according to claim 10 wherein the depot is connected to one or more servers through a local area network or a frame relay network, comprising further steps of:

performing the TCP connection sessions in their appropriate format through the local area network or the frame relay network.

15. The method according to claim 10, wherein the depot is connected to one or more servers by way of a plurality of proxy servers, comprising further steps of:

switching, at the depot, the TCP connection sessions among the proxy servers which in turn send TCP connection sessions to the server; and

caching information resources contained in the TCP connection sessions at the proxy servers.

16. The method according to claim 15, comprising further steps of:

monitoring the status of all the proxy servers; and

for each TCP connection session, selecting one proxy server among all the proxy servers based on their monitored status.

17. The method according to claim 16 wherein there are plurality of depots, comprising further steps of:

monitoring parameters indicative of the operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

18. The method according to claim 14 wherein there are a plurality of depots, comprising further steps of:

monitoring parameters indicative of the operability of all the depots; and

upon detection of inoperability of one depot, continuing all the functions of the failed depot at any of the remaining depots.

19. In a client-server environment of a telecommunications network, where information resources are replicated among a plurality of servers and a plurality of clients have access to any one or more of the servers for desired information resources through an HTTP transaction using at least two layers of protocols, one stateless protocol layer upon another stateful protocol layer, the HTTP transaction running on the stateless protocol layer of HTTP protocol and being made up of one or more TCP connection sessions which run on TCP/IP layer, a depot for transferring one or more TCP connection sessions between a client and among a plurality of servers, comprising:

a primary map and a secondary map for storing information about the TCP connection sessions and parameters, the information at the primary map being deleted at the close of sessions and the information at the secondary map being maintained for two maximum segment lifetimes (MSL);

a packet analysis block for analysing and transparently forwarding TCP/IP packets of one or more TCP connection sessions of the HTTP transaction between the client and one of the servers;

a session management block for switching one or more TCP connection sessions of the HTTP transaction among the plurality of servers according to a preset criteria including the TCP/IP states of each TCP connection session; and

a protocol state tracking block for monitoring the TCP/IP states of the TCP connection sessions of the HTTP transaction.

20. The depot according to claim 19, further comprising:

a map which contains data concerning either one or both of load and operability of the plurality of servers; and

the session management block for switching one or more TCP connection sessions of the HTTP transaction among the servers based on the data stored in the map.

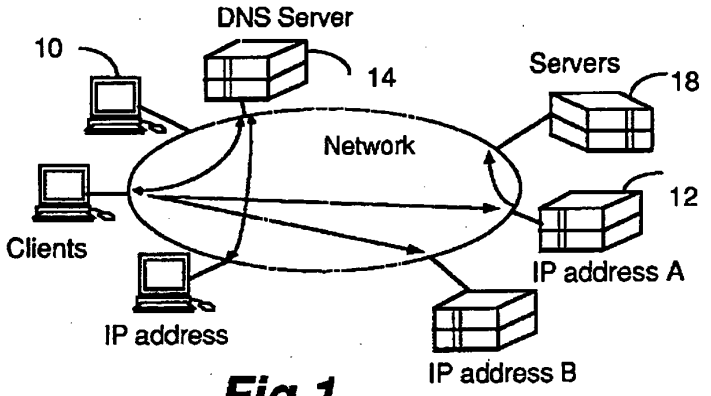


Fig 1

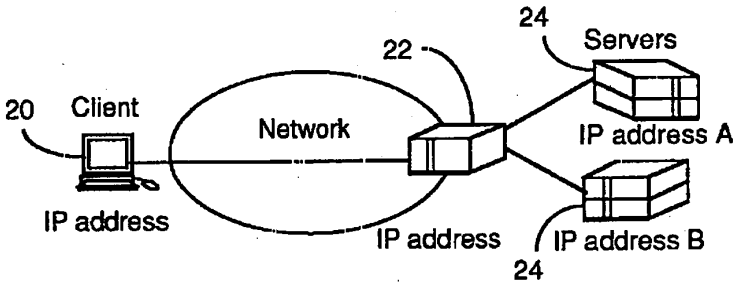


Fig 2

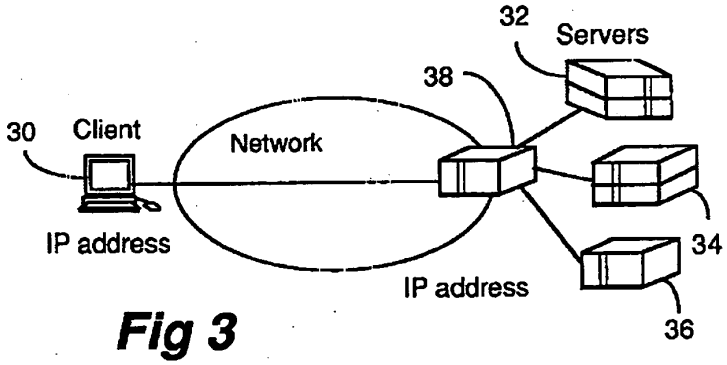


Fig 3

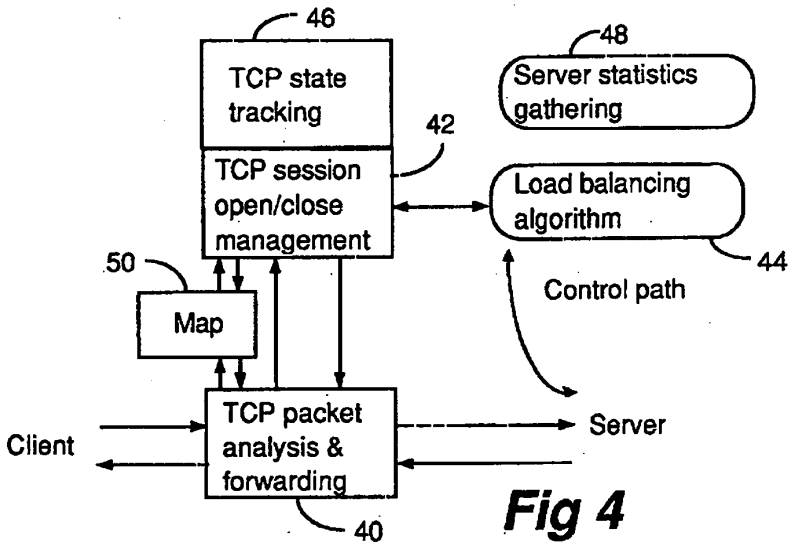

46
TCP state tracking
48
Server statistics gathering
42
TCP session open/close management
Load balancing algorithm
44
50
Map
Control path
Client
TCP packet analysis & forwarding
Server
40

Fig 4

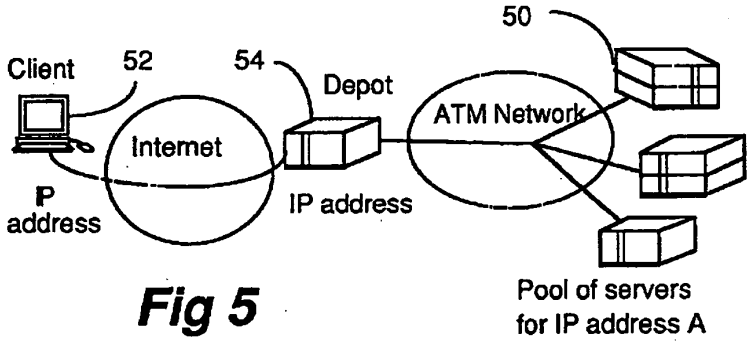

50
Client
52
54
Depot
ATM Network
Internet
IP address
IP address
Pool of servers for IP address A

Fig 5

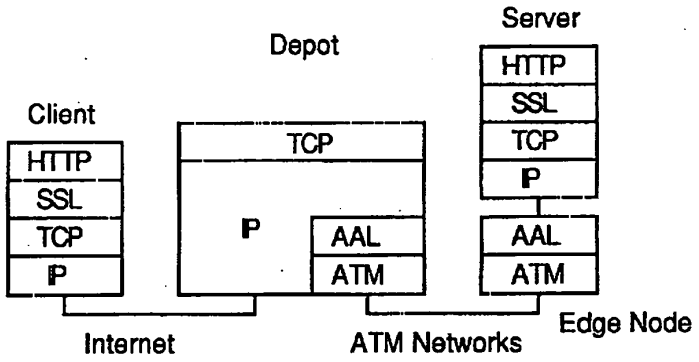


Fig 6

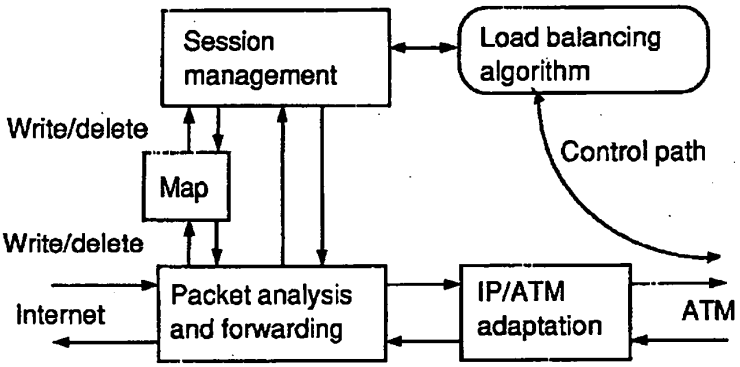


Fig 7

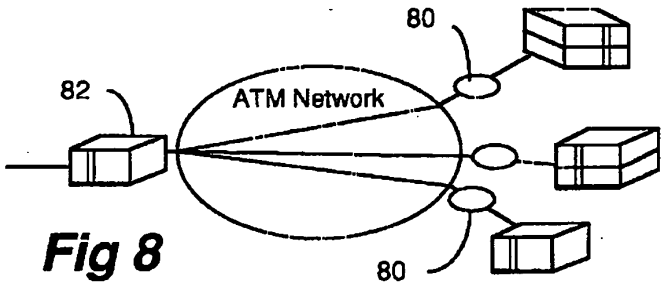


Fig 8

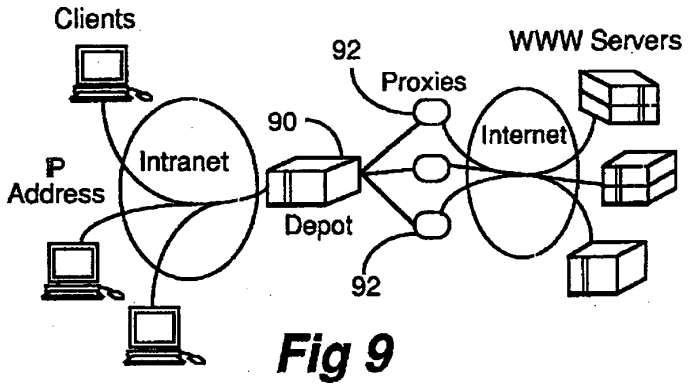

Clients
92
Proxies
WWW Servers
90
Internet
IP
Address
Intranet
Depot
92

Fig 9

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 865 180 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
16.09.1998 Bulletin 1998/38

(51) Int Cl.[6]: **H04L 12/56**, H04L 29/06, H04L 29/12

(21) Application number: **98301577.7**

(22) Date of filing: **03.03.1998**

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **14.03.1997 US 818989**

(71) Applicant: **LUCENT TECHNOLOGIES INC.**
Murray Hill, New Jersey 07974-0636 (US)

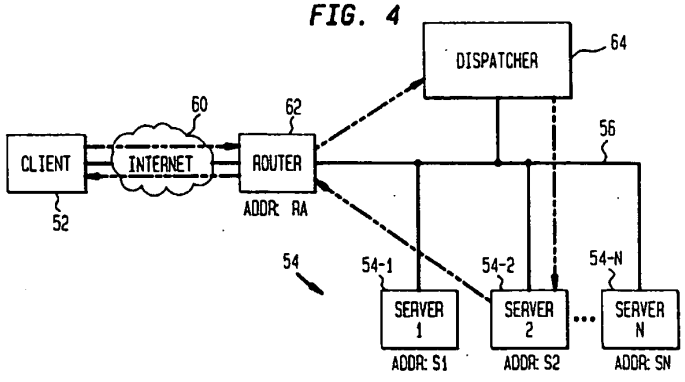
(72) Inventors:
- **Chung, Pi-Yu**
  Berkeley Heights, New Jersey 07922 (US)
- **Damani, Om P.**
  Austin, Texas 78705 (US)
- **Huang, Yennun**
  Bridgewater, New Jersey 08807 (US)
- **Kintala, Chandra M.**
  Warren, New Jersey 07059 (US)
- **Wang, Yi-Min**
  Berkeley Heights, New Jersey 07922 (US)

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr.
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(54) Load distribution among servers in a TCP/IP network

(57) Methods and apparatus for hosting a network service on a cluster of servers, each including a primary and a secondary Internet Protocol (IP) address. A common cluster address is assigned as the secondary address to each of the servers in the cluster. The cluster address may be assigned in UNIX-based servers using an *ifconfig alias* option, and may be a ghost IP address that is not used as a primary address by any server in the cluster. Client requests directed to the cluster address are dispatched such that only one of the servers of the cluster responds to a given client request. The dispatching may use a routing-based technique, in which all client requests directed to the cluster address are routed to a dispatcher connected to the local network of the server cluster. The dispatcher then applies a hash function to the client IP address in order to select one of the servers to process the request. The dispatching may alternatively use a broadcast-based technique, in which a router broadcasts client requests having the cluster address to all of the servers of the cluster over a local network. The servers then each provide a filtering routine, which may involve comparing a server identifier with a hash value generated from a client address, in order to ensure that only one server responds to each request broadcast by the router.

FIG. 4



EP 0 865 180 A2

Printed by Jouve, 75001 PARIS (FR)

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 824 236 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
18.02.1998 Bulletin 1998/08

(51) Int Cl.[6]: **G06F 9/46**

(21) Application number: **97304739.2**

(22) Date of filing: **01.07.1997**

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV RO SI

(30) Priority: **08.08.1996 GB 9616649**

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**
Armonk, NY 10504 (US)

(72) Inventors:
- **Baird, Richard**
  Markham, Ontario L3R 8V3 (CA)
- **Ogilvie, Alan**
  Southampton, Hampshire SO52 9NA (GB)

(74) Representative: **Davies, Simon Robert**
I B M
UK Intellectual Property Department
Hursley Park
Winchester, Hampshire SO21 2JN (GB)

(54) Computer system

(57) A computer workstation 1 is attached to a network such as the Internet 34. Browser software 210 running on the computer workstation can access and retrieve pages of data from over the network, with the pages of data then being displayed at the computer workstation. A page of data can include a reference to code which the browser can also retrieve from over the network and then execute on the computer workstation. The reference to the code also has associated with it a reference to another file, again accessible over the network, which contains information pertaining to the code, such as the size of the code. The browser accesses the information in this file, prior to retrieving or executing the code. This information can then help the user at the computer workstation determine whether or not to retrieve the code.

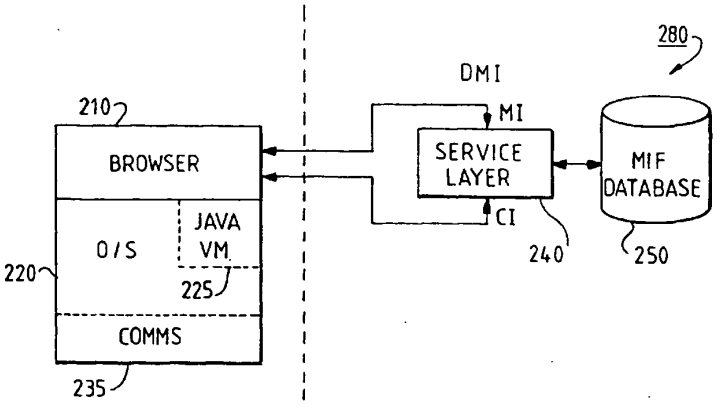


FIG. 2

EP 0 824 236 A2

Printed by Jouve, 75001 PARIS (FR)

Description

The present invention relates to a computer system, and to a method of operating such a computer system attached to a network, such that the computer system can access data and code over the network.

In the past couple of years there has been an explosive growth in the Internet, and in particular of the WorldWide Web (WWW), which is one of the facilities provided on top of the Internet. The WWW comprises many pages or files of information, distributed across many different servers. Each page is identified by a Universal Resource Locator (URL). The URL denotes both the server machine, and the particular file or page on that machine. There may be many pages or URLs resident on a single server.

In order to use the WWW, a client runs a piece of software known as a Web browser, such as WebExplorer (provided as part of the OS/2 operating system from IBM Corporation), or the Navigator program available from Netscape Communications Corporation. The client interacts with the browser to select a particular URL, which in turn causes the browser to send a request for that URL or page to the server identified in the URL. Typically the server responds to the request by retrieving the requested page, and transmitting the data for that page back to the requesting client (the client server interaction is performed in accordance with the hypertext transport protocol ("http")). This page is then displayed to the user on the client screen. The client may also cause the server to launch an application, for example to search for WWW pages relating to particular topics.

Most WWW pages are formatted in accordance with a language known as html (hypertext mark-up language). Thus a typical page includes text together with embedded formatting commands, referred to as tags, which can be used to control the font size, the font style (for example, whether italic or bold), how to lay-out the text, and so on. A Web browser parses the html script in order to display the text in accordance with the specified format. In addition, an html page can also contain a reference, in terms of another URL, to a piece of multimedia data, for example, an image, a video segment, or an audio file. A Web browser responds to such a reference by retrieving and displaying or playing the data. Alternatively, such multimedia data may form its own WWW page, without any surrounding html text.

Most WWW pages also contain one or more references to other WWW pages, which need not be on the same server as the original page. Such references may generally be activated by the user selecting particular locations on the screen, typically by (double) clicking a mouse control button. These references or locations are known as hyperlinks, and are typically flagged by the browser in a particular manner (for example, any text associated with a hyperlink may be in a different colour). If a user selects the hyperlink, then the referenced page is retrieved and replaces the currently displayed page.

Further information about html and the WWW can be found in "World Wide Web and HTML" by Douglas McArthur, p18-26 in Dr Dobbs Journal, December 1994, and in "The HTML SourceBook" by Ian Graham, (John Wiley, New York, 1995).

As so far described, and broadly speaking as currently implemented, the WWW suffers from the disadvantage that pages downloaded from a server to a client are essentially passive, in other words, they do not contain code which is executed at the client machine. One implication of this is that the server cannot offload onto the client any of the processing associated with the interaction between the client and the server. Thus if the client is completing a form with their telephone number for example, then any formal checks such as to the number of digits in the telephone number must be performed at the server. This results firstly in a heavier processing burden at the server, and secondly in time-consuming extra communications between the server and client should there be any mistakes to correct. Moreover, the inability of the server to download code for execution at the client is a significant limitation on the type of applications that can be created to exploit the WWW.

Recent developments, based particularly on the Java technology from Sun Microsystems Inc, have sought to overcome the above difficulties. The Java technology comprises primarily (i) a new programming language, somewhat similar to C and C++, and (ii) a virtual machine. Essentially, programs written in the Java programming language can be compiled into byte code form, and then interpreted at runtime on the Java virtual machine. The Java virtual machine converts the byte codes into instructions that can be executed by the underlying physical machine.

Programs written using Java can be downloaded over the WWW in the form of byte codes for execution on a Java virtual machine at the client. Such programs are known as "applets". The use of the Java technology for downloading code over the WWW has two major benefits. Firstly, an applet can be platform independent, if we assume that each client has a copy of the Java virtual machine (the virtual machine at the client's system is typically incorporated either into the operating system, or into the Web browser itself). In other words, there is no need for a server to have different versions of the code for downloading to clients according to their respective operating systems and machines. Therefore, only a single version of the relevant code needs to be written and maintained, which makes life much simpler for software developers. Secondly, because the applet executes on a virtual machine, rather than a physical machine, security is greatly improved. Thus, when downloading code over the network, there is always a risk that it will include some malicious code (accidentally or otherwise) that may damage data or programs stored at the client. The virtual machine however can monitor the operation of the applet, and so detect and prevent such malicious activity.

It will be noted that the concept of downloading software from a server to a client in the form of byte codes for execution on a virtual machine was also known independently of the Java technology, see for example US 5347632.

In order to invoke a Java applet, a Web page of html text contains an <APPLET> tag, which identifies the URL containing the applet. A browser responds to this tag by retrieving and running the applet. Also defined is a <PARAM> tag, which is contained within a pair of corresponding <APPLET> and </APPLET> tags, and which can be used to specify parameters that are passed to the applet at run-time. (Note that the APPLET and PARAM tags are not formally incorporated into the html standard, but are nevertheless recognised by many Web browsers). Further information about the Java technology and applets can be found in "Teach Yourself Java in 21 Days" by Laura Lemay and Charles Perkins (Sams.net Publishing, Indianapolis, USA, 1996).

One of the difficulties with the above approach is that the amount of information that can be contained within the APPLET tag is limited. This can lead to problems if the browser has to decide whether or not to download a particular applet referenced within an html page. For example, the browser may not wish to download an applet which is particularly large (and hence slow to transport across the network), or which requires a more recent version of the Java virtual machine than the one currently installed at the client.

Accordingly, the invention provides a method of operating a computer attached to a network, having accessible at least one page of data that includes a reference to code that can be retrieved over the network for execution at said computer, characterised by the steps of: locating, associated with said reference, an identifier to information about said code; and retrieving the information about the code.

Thus the computer can utilise the identifier to locate and retrieve information concerning the code. This information may then be used in order to determine whether or not to run said code. For example, the information preferably states the size of the code, and it may be decided that the code is too large to be downloaded within a reasonable time; another possibility is that the information specifies that execution of the code requires hardware support features, such as a video capture facility, which may not be installed at the client workstation. Such a determination may be explicitly surfaced to the user, or may be performed automatically based on predetermined criteria (or a combination of both automatic and user-prompted decision making could be adopted).

Other possible uses for the information about the code might include for audit purposes, perhaps in circumstances where there is some royalty payment attached to use of the code, or for simple informational purposes, such as providing access to a copyright notice and licensing terms, and so on. Another alternative is to use the information to assist in invoking the code itself. For example, if the code is encrypted, then the information may specify the algorithm needed to decrypt the code before it can be invoked.

In the preferred embodiment, once the information about the code has been retrieved and stored and/or processed as desired, the code itself is then retrieved over the network. Alternatively, the code might be retrieved prior to retrieval of the information about the code.

In the preferred embodiment, the at least one page of data is in html format, the reference comprises an APPLET tag, and the identifier comprises one or more attributes within said APPLET tag. Preferably said reference and said identifiers are in the form of Universal Resource Locators (URLs), with each reference and identifier being specified in two parts, the first specifying server and directory, the second the actual file name. In this preferred embodiment, the applet is downloaded in Java byte code form, for execution on a Java virtual machine at the client's workstation. This represents the typical situation where the computer is running a Web browser, and said at least one page of data, said code reference and said information are all accessible over the Internet.

It will be appreciated that since the specific APPLET tag is particularly associated with the Java language, it may be replaced in some future standard with a more generalised tag (ie one that is regarded as more language-neutral), although of course this would not affect the principle of the present invention. It will also be appreciated that in some circumstances, said at least one page of data may contain little or nothing apart from the reference to the code, and its associated identifier of the information (in other words, the primary or sole purpose of the page may be to access the code).

In a preferred embodiment, the information conforms to the Management Information Format (MIF). This is a well-known standard, and so is likely to be understood by the client trying to access the code. However, the information could be stored in any format which the browser can understand and access. The MIF information is copied over the network and then preferably loaded into the MIF database at the computer, where it can be interrogated by the browser (or simply stored for audit purposes). Alternatively, the MIF information may already be stored in a remote MIF database, so that information about the code can be retrieved by remotely interrogating said database over the network. This saves having to transmit the whole MIF file over the network.

Thus the invention provides a mechanism for making available information about code which can be accessed over the network. An alternative to the approach provided by the invention might be for example to have information relating to code present in an html page itself, for example as additional attributes in the <APPLET> tag. However, this would not only lead to the tags becoming potentially very large, but also each page which referenced a particular applet would then have to contain a copy of this information. This is not only more expensive in terms of storage requirements,

but also extremely difficult to maintain from a practical point of view, since changes to the applet have to be reflected at multiple locations (indeed the owner of the code will generally be unaware of all the sites which reference that code). Another alternative might be for the information about code to be identified within the code itself, so that on execution the code initially brings up the relevant information. This strategy has the drawback that it requires the code to always be at least partly executed. Thus it can be seen that the approach adopted by the present invention represents the most powerful, flexible, and efficient way of making available information about code which can be accessed over the network.

The invention also provides a computer system, attachable to a network to access at least one page of data that includes a reference to code that can be retrieved over the network for execution at said computer, characterised by: means for locating, associated with said reference, an identifier to information about said code; and means for retrieving the information about the code.

The invention further provides a computer program product for execution by a computer system attached to a network to access at least one page of data that includes a reference to code that can be retrieved over the network for execution at said computer system, said computer program product including: means for locating, associated with said reference, an identifier to information about said code; and means for retrieving the information about the code.

The computer program product typically resides on a diskette or a CD-ROM, or alternatively may be stored on a tape or hard disk of a server machine for downloading across the network, and in use is copied into the working memory (typically RAM) of the computer system.

The invention further provides a Web browser for accessing at least one page of data on the World Wide Web that includes a reference to applet code that can be retrieved over the Internet for execution at said computer, characterised by: means for locating, associated with said reference, an identifier to information about said applet code; and means for retrieving the information about the code.

The invention further provides a method of operating a server attached to a network, said method including steps of: storing at least one page of data that includes a reference to code that can be retrieved over the network, and an identifier associated with said reference to indicate the location of information stored about said code; receiving a request over the network for said at least one page of data; and responsive to said request, transmitting said at least one page of data.

A preferred embodiment of the invention will now be described in detail by way of example only with reference to the following drawings:

Figure 1 is a simplified schematic diagram of a computer system;

Figure 2 is a simplified schematic diagram of the main software components on the computer system of Figure 1;

Figure 3 is a flowchart illustrating processing according to a first embodiment of the invention; and

Figure 4 is a flowchart illustrating processing according to a second embodiment of the invention.

Figure 1 is a simplified schematic diagram of a computer system 1, including a system unit 10, a display screen 12, a keyboard 14 and a mouse 16. The system unit includes a microprocessor 22, semi-conductor memory (ROM/RAM) 24, a hard disk drive 28, and a bus over which data is transferred 26. A network adapter card 30 is provided, which, in conjunction with link 32, enables the computer workstation to exchange data with the Internet 34. The computer is essentially controlled by the microprocessor, based on software instructions loaded into ROM/RAM, typically from the hard disk drive, or from the network. The system of Figure 1 may be any conventional workstation, such as a Personal Computer machine or a RISC System/6000 computer, both available from IBM Corporation. Such systems are very well-known in the art, and the skilled person is aware of great variety of possible system configurations: for example, the system may include a CD-ROM, have multiple network adapter cards, use a telephone modem instead of a network adapter card for the network connection, include audio systems, and so on. Such variations are not pertinent to an understanding of the present invention.

Figure 2 illustrates the software components running on system 1, including the operating system 220 and Web browser 210 (note that in some cases the browser may be incorporated into the operating system). As previously discussed, the browser is responsible for accessing material over the Internet, and displaying or otherwise processing the retrieved material as appropriate. One important feature of the browser is its ability to understand html commands. The operating system 220, such as AIX or OS/2, both available from IBM Corporation is responsible for general system management of the workstation. Included within the operating system is communication software 235 which allows the workstation to interact with the Internet 34. As is well known in the art, such communications are based on the TCP/IP protocol. The operating system also includes a Java virtual machine 225, which can interpret Java byte codes in order to run Java applications. Note that in some cases the Java virtual machine may be provided as part of the

browser. Also running in one embodiment on system 1 is Desktop Management Interface (DMI) software 280 (which also may be incorporated into the operating system 220). This will be described in more detail later.

System 1 and its software can be used to access and process WWW pages, including Java applets, in known fashion, using the http protocol. An example of the html text in a Web page which references a Java applet is as follows:

```
<APPLET CODEBASE = "http://ncc.hursley.ibm.com/javainfo/classes"
      CODE = "Map.class"
      WIDTH = 560
      HEIGHT = 560
>
<param name = IMAGE value = "graphics/cup-bun.jpg">
<param name = MAP value = "java.map">
<param name = TYPE value = "NCSA" >
</APPLET>
```

Example 1

The browser interprets this html text as follows. In response to the <APPLET> and </APPLET> tags, the browser retrieves the code which is located in file "Map.class" at the server and directory specified by "http://ncc.hursley.ibm.com/javainfo/classes". (As is well-known in html, tags of the form "</xxx>" mark the end of information or data associated with the preceding corresponding <xxx> tag). This processing is essentially the same as accessing any WWW page from a URL. The browser then creates an applet context object. Associated with this are the values of HEIGHT and WIDTH (known as attributes), which define the amount of screen which the applet needs for execution, and also certain parameter values. These parameter values are listed in their own <PARAM> tags which specify the name and value of the tag, and can be retrieved into the applet itself by the use of the Java getParameter method. Thus for example TYPE is the name of a string variable in the Map.class applet, which is to be assigned the value "NCSA". Finally, the applet is executed by creating an applet object using the parameter data previously stored in the applet context object and then invoking the applet object using the Java **Init()** instruction.

As mentioned previously, some difficulties can arise with this known approach, for example, the Java applet may require a more recent version of the Java virtual machine than is installed at the client. In this situation, if the client tries to download and run the applet, he or she may discover that operation of the applet is limited or impossible.

The present invention overcomes this problem by expanding the information presented to the browser relating to the applet. Thus, in accordance with the present invention, the html text of Example 1 would be modified as follows:

```
<APPLET CODEBASE = "http://ncc.hursley.ibm.com/javainfo/classes"
      CODE = "Map.class"
      MIFBASE= "http://ncc.hursley.ibm.com/javainfo/mifs"
      MIF="Map.mif"
      WIDTH = 560
      HEIGHT = 560
>
<param name = IMAGE value = "graphics/cup-bun.jpg">
<param name = MAP value = "java.map">
<param name = TYPE value = "NCSA" >
</APPLET>
```

Example 2

It will be noted that Example 2 is the same as Example 1, except for the addition of two new attributes into the <APPLET> tag. The two attributes, MIFBASE and MIF, are analogous to the fields CODEBASE and CODE, in that they specify a URL comprising a server/directory and file respectively. However, whereas CODEBASE and CODE locate the actual code for the Java applet, MIFBASE and MIF locate a file containing information describing the applet. This file may reside on the same server and possibly in the same directory as the applet code itself, or may reside on a completely different server.

In the preferred embodiment, the file containing the information describing the applet is in a particular format known as the Management Information Format (MIF). This is a standard format developed by the DeskTop Management TaskForce (DMTF) industry consortium as part of the Desktop Management Interface (DMI). A short overview of DMTF and MIF is given in "The DMI Waiting Game" by Daniel Begun, in the Computer Shopper, n8, v15, p608, August 1995, and in "The End of the Balancing Act" by Barbara Gengler, in Internetwork, n8, v6, p47, August 1995. More information about MIF can be obtained from the DMTF at Hillsboro, Oregon, USA, or from their Web site, see in particular:

http://wwww.dmtf.org/general_info/factsheet.html, and
http://wwww.dmtf.org/techlinks/white_papers.html

Briefly, a MIF file, which is in the form of plain text, comprises multiple groups. Certain standard groups have been defined, some of which are mandatory, but user-defined groups are also permitted. Within each group is a list of named attributes and their values. For example, one required group is "ComponentID", which has the attributes of Manufacturer, Product, Version, Serial Number, Installation (date and time), and Verify (whether the corresponding component actually exists and is working). Referring back to Figure 2, a system supporting the DMI 280 has a MIF database 250, and a service layer 240 (these may be provided as part of the machine operating system). Manageable products (such as software applications) add or update their respective MIF files into or in the MIF database via a Component Interface (CI) of the service layer, whilst management applications can query the MIF database for information about the manageable applications via the Management Interface (MI) of the service layer. Samples of software MIF files plus further descriptive material can be obtained from the following Web site:

http://www-uk.hpl.hp.com/people/arp/dmtf/ver2.htm

The operation of the invention in one simple embodiment is shown schematically in Figure 3. Initially, the browser software detects the presence of an <APPLET> tag (step 310), representing a reference to an applet to be downloaded and executed. Next the browser looks for the MIFBASE and MIF attributes associated with the referenced applet (step 320). Assuming that these are located, the browser then accesses and retrieves the specified MIF file (steps 330 and 340). The MIF file is then displayed to the user in a window (step 350), allowing the user to scroll through the MIF information, and the user is prompted with a dialog box asking whether or not to continue with the applet (step 360). Assuming that the MIF information is acceptable, the user can then download and invoke the applet (steps 370 and 380). However, if the MIF file reveals some problem, such as the applet being extremely large, or requiring a more recent version of the Java virtual machine than is installed at the client, then the user has the option of effectively aborting, thereby avoiding needlessly downloading or trying to run the applet.

Returning to step 320, if the MIF and MIFBASE attributes are not located, then an attempt is made to download and invoke the applet directly. Note that this is also the path taken by a browser which does not recognise the MIF and MIFBASE attributes, which would be the case with a prior art browser.

Although the approach shown in Figure 3 is relatively simple, it has the drawback that the user has to scan through the MIF file itself. A more complex processing approach is depicted in the embodiment of Figure 4 (where we assume that the APPLET tag does contain the MIF/MIFBASE attributes). In this case, once the MIF information has been retrieved (step 405) from the server, it is loaded into the MIF database on the computer system (410) via the component interface (CI) of the service layer. Then, the management interface (MI) of the service layer is used to read out the MIF information from the MIF file (step 415). This interaction between the browser and the DMI is shown schematically in Figure 2.

The browser tests the MIF information against predetermined criteria (step 420), such as to determine whether the version number of the target operating system (the Java virtual machine) is compatible with the version installed at the client (note that the browser can use the DMI to retrieve only selected information from the MIF file, ignoring attributes that are not needed for this testing). Unless the predetermined criteria are satisfied, processing of the applet terminates, without the applet ever having been downloaded. There may be some criteria which are always or conditionally surfaced to the user for a decision; for example, if the applet is larger than a certain size, the user may be asked to confirm whether or not to continue (step 425), bearing in mind the length of time that may be required in order to download the applet. Assuming that the predetermined criteria are satisfied (with any necessary user confirmation), the applet is then downloaded (step 430). At this stage, further testing can be performed (step 435), typically for security purposes. For example, it can be verified that the applet is the correct size, and checks made that certain bytes specified in the MIF file properly match the content of the applet itself. Assuming that the predetermined criteria are once more

satisfied (again subject to any necessary user confirmation at step 440), then the applet can now be invoked (step 445). Note that if the browser subsequently deletes the applet from the client system, then generally it should update the MIF database accordingly.

It will be appreciated that a great many variations are possible on the processing illustrated in Figures 3 and 4. For example, the MIF file may simply be loaded into the MIF database (step 410) for audit purposes, without performing any checking against predetermined criteria, or asking for user verification. Typically this would be done after the applet has been downloaded (step 430), with the browser then attempting to invoke the applet (step 445), irrespective of the content of the MIF file. Alternatively, some browsers may do the testing against predetermined criteria all automatically, without any of the testing being surfaced to the user. Another possibility is that the MIF file could contain information needed to execute the applet. For example, it has been contemplated that applets could be downloaded in encrypted form, to reduce the possibility of tampering, in which case the MIF could be used to specify the algorithm needed to decrypt the applet.

It is optional whether the MIF file is loaded into the MIF database, and if so, whether the browser obtains the MIF information directly from the MIF file itself, or by using the MIF database. Loading the MIF file into the MIF database may be useful for audit purposes, and also allows the browser to exploit the MI for retrieving MIF information from the MIF file. This avoids the need for the browser to have to parse the MIF file to retrieve the values of the various attributes. However, if a large number of applets are downloaded onto the client machine, and have only a transitory existence there before being executed and deleted, then in some circumstances it may be preferable to obtain the MIF information directly from the MIF file itself, rather than having the overhead associated with interacting with the DMI.

Note that some of the MIF attributes discussed, such as size of the code, and version number of the target operating system, are already present in standard DMI groups, whilst others, such as type of encryption used, would need to be placed in user-defined groups. There are no firm limits on the type of information that can be contained within a MIF, although clearly it is desirable to have a set of predefined attributes, preferably as part of the DMI standard, which most or all browsers can interpret. It is also desirable that if a browser does not recognise particular attributes in this set, then it is still able to invoke the associated applet where possible (this is somewhat analogous to the situation where a browser does not recognise a particular html tag, but can still generally display the information on a page, even if not in exactly the desired format).

In the above embodiments, the MIF file is downloaded in its entirety for analysis at the client. However, an alternative approach is to retrieve the information concerning the applet by interrogating the associated MIF file at a remote location. In this case it is most likely that the MIF file is incorporated into a MIF database, so that the attributes MIF and MIFBASE would typically be used to locate the DMI through which access to the MIF database is provided. This arrangement is essentially the same as that shown in Figure 2, except that now the DMI 280 is located on a different machine from operating system 220 and browser 210.

There are several possible mechanisms for obtaining this remote access to the DMI. For example, it is intended that DMI itself should directly support such remote access via the MI, although a full implementation of this is not directly available yet. In this case, the processing might be similar to that shown in Figure 4, except that step 410 would be omitted, and the retrieval of the MIF information in step 415 would be performed via the remote access facility. Alternatively, the browser could use the well-known Simple Network Management Protocol (SNMP) to remotely interrogate the DMI, since DMI already supports this mechanism (the SNMP is described in "Internetworking with TCP/IP, Volume 2", by Douglas E Comer and David L Stevens, Prentice Hall, New Jersey, USA, 1991). Another possibility, which would avoid the need for the browser to include SNMP support, would be to have MIF and MIFBASE identify an intermediary program on the same machine as the MIF database. This intermediary would then receive requests in the http protocol from the browser for MIF information, with the requests specifying the particular desired attribute(s), submit these requests to the MI of the DMI, and then return the response to the browser, again in the http protocol.

The use of the MIF information directly at the remote location may be more efficient than initially copying the MIF information to the client, particularly if the browser is only interested in relatively few of the MIF attributes, since this avoids the need for the whole MIF file to be transmitted across the network. Of course, with this approach the browser would have to separately record any desired audit information.

It will be appreciated that although in the preceding description the identifier for the applet information has been broken down into two components designated MIFBASE and MIF, the identifier could be represented in any suitable form and under any appropriate name. Further, the information about the code need not be a MIF file, but could be in a different format, such as the SNMP management information base (MIB) format, or any other suitable format defined or specified by the user. Also, whilst the identifier in Example 2 is located within the APPLET tag itself, one could introduce a separate tag located between the <APPLET> and </APPLET> tags (ie analogous to the <PARAM> tag) in which to store the identifier. Of course, with all the above considerations, it is desirable to standardise on one particular option, which can then be supported by most or all browsers.

Further, whilst in the embodiments described above, the code for reviewing and interpreting the MIF file has been incorporated into the browser, this code could also form a separate program or object residing at the client, and known

to the browser. Thus on retrieving the MIF file, the browser would pass the MIF file to the MIF interpreter program, which would then analyze the file in the desired manner. For example, the interpreter program could simply open a window to display the MIF file, and possibly also ask the user whether to continue, corresponding to steps 350 and 360 in Figure 3; or simply add the MIF file to the local MIF database for auditing purposes, or test the MIF file against predetermined criteria. Where appropriate, the MIF interpreter program would return an appropriate response to the browser, indicating whether or not it should proceed with execution of the applet.

Alternatively, the MIF interpreter program could also be referenced by the Java virtual machine itself, which when invoked to run an applet, initially runs the MIF interpreter program, for example to record use of the applet for auditing purposes, or to ensure that the installed version of the virtual machine is compatible with that required for running the applet. Thus the MIF interpreter would act as a form of preprocessor in the virtual machine, in a similar manner (and immediately prior to) the byte code validator which currently initially verifies byte codes passed to the virtual machine for execution. In this approach the virtual machine preferably contains a hook to allow third party MIF interpreter programs to be bolted on to the virtual machine, although it would also be possible to incorporate that MIF interpreter function into the virtual machine itself.

Additionally, although the preferred embodiment has been described primarily with reference to an otherwise conventional computer workstation attached to the Internet, it will be appreciated that the invention is not so limited. For example, the computer workstation might be replaced by a so-called network computer, or any other hardware which is capable of receiving code over a network, such as a set-top box, a television Internet access device, and so on. Likewise, the relevant network may not be the Internet, but might be a corporate intranet (which typically uses Internet protocols such as html, but which is detached from the Internet itself), or any other suitable network. Further, whilst the code to be downloaded in the preferred embodiment comprises Java byte codes, this code might be in any appropriate language, and might be in a form either to be interpreted, for example by a virtual machine, as with Java, or to be directly executed on the hardware of the client machine. In this case the MIF file could be used to verify that the target operating system of the code to be downloaded matches that installed at the client machine.

Claims

1. A method of operating a computer attached to a network, having accessible at least one page of data that includes a reference to code that can be retrieved over the network for execution at said computer, characterised by the steps of:

   locating, associated with said reference, an identifier to information about said code; and

   retrieving the information about the code.

2. The method of claim 1, further comprising the step of utilising said information in order to determine whether or not to run said code.

3. The method of claim 2, further comprising the step of presenting at least part of said information to the user of said computer in order to allow said user to determine whether or not to run said code.

4. The method of claim 2 or 3, further comprising the step of comparing at least part of said information against one or more predetermined criteria to determine whether or not to run said code.

5. The method of any preceding claim, wherein said at least one page of data is in html format.

6. The method of claim 5, wherein said reference comprises an APPLET tag.

7. The method of claim 6, wherein said identifier comprises one or more attributes within said APPLET tag.

8. The method of claim 7, wherein said identifier comprises a first attribute specifying the server and directory where said information is located, and a second attribute specifying the name of a file containing the information.

9. The method of any preceding claim, wherein said computer runs a Web browser to access said at least one page of data, and said reference is in the form of a Universal Resource Locator (URL).

10. The method of any preceding claim, wherein the information is stored in a file, and said step of retrieving the

information about the code comprises copying the file onto said computer.

11. The method of claim 10, wherein the information is stored in a form conforming to the Management Information Format (MIF), and is loaded into the MIF database at the computer.

12. The method of claim 11, further comprising interrogating said MIF database to access the loaded MIF information.

13. The method of any of claims 1 to 9, wherein the information is stored in a remote database, and said step of retrieving the information comprises remotely interrogating said database over the network.

14. The method of any preceding claim, wherein said information includes the size of said code.

15. The method of any preceding claim, further comprising the steps of using said reference to retrieve the code over the network and executing the retrieved code at the computer.

16. A computer system, attachable to a network to access at least one page of data that includes a reference to code that can be retrieved over the network for execution at said computer, characterised by:

means for locating, associated with said reference, an identifier to information about said code; and

means for retrieving the information about the code.

17. A Web browser for accessing at least one page of data on the World Wide Web that includes a reference to applet code that can be retrieved over the Internet for execution at said computer, characterised by:

means for locating, associated with said reference, an identifier to information about said applet code; and

means for retrieving the information about the code.

18. The browser of claim 17, wherein said reference comprises an APPLET tag, and said identifier comprises one or more attributes within said APPLET tag.

19. A method of operating a server attached to a network, said method including steps of:

storing at least one page of data that includes a reference to code that can be retrieved over the network, and an identifier associated with said reference to indicate the location of information stored about said code;

receiving a request over the network for said at least one page of data;

and responsive to said request, transmitting said at least one page of data.

20. The method of claim 19, wherein said reference comprises an APPLET tag, and said identifier comprises one or more attributes within said APPLET tag.

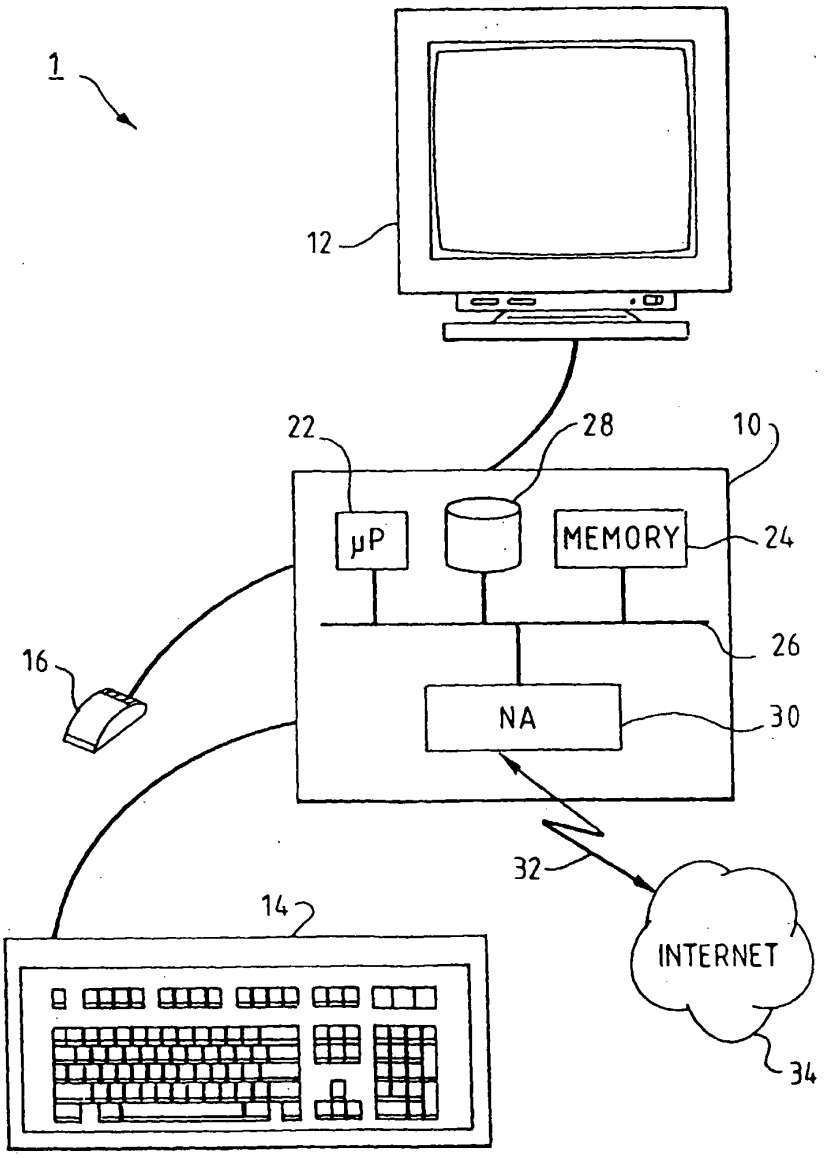


FIG. 1

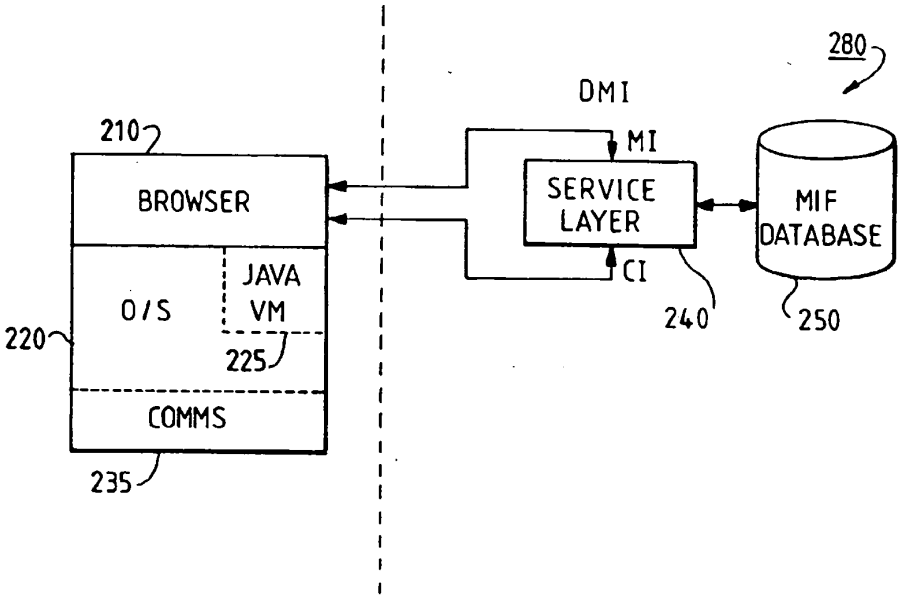

280
DMI
MI
210
BROWSER
SERVICE LAYER
MIF DATABASE
CI
O/S
JAVA VM
220
225
240
250
COMMS
235

FIG. 2

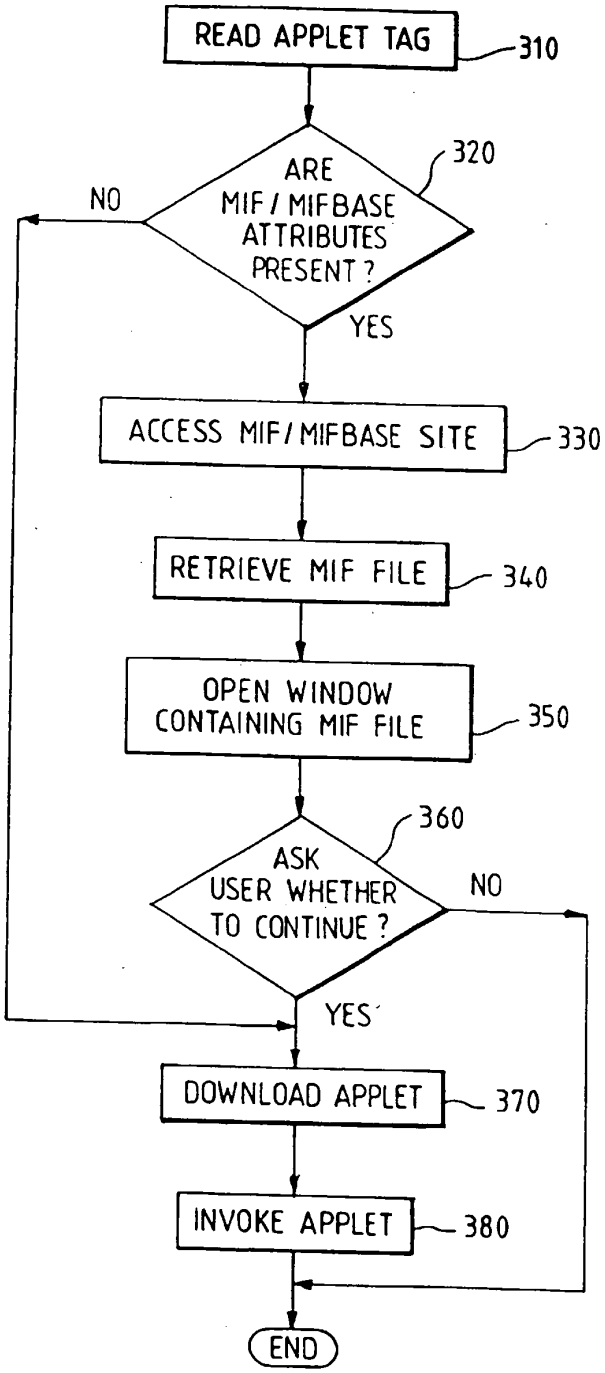

READ APPLET TAG
310
ARE MIF / MIFBASE ATTRIBUTES PRESENT ?
320
NO
YES
ACCESS MIF/MIFBASE SITE
330
RETRIEVE MIF FILE
340
OPEN WINDOW CONTAINING MIF FILE
350
ASK USER WHETHER TO CONTINUE ?
360
NO
YES
DOWNLOAD APPLET
370
INVOKE APPLET
380
END

FIG. 3

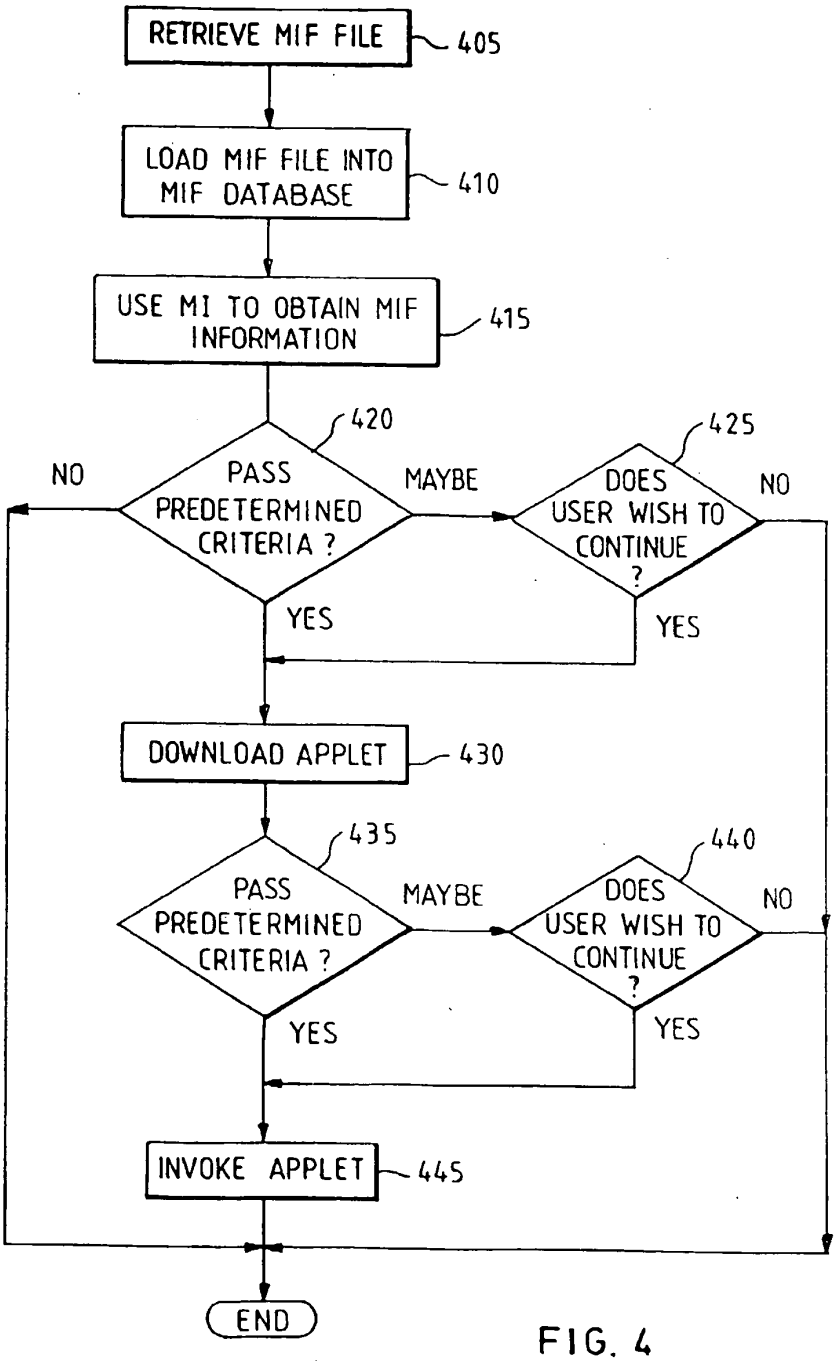

RETRIEVE MIF FILE
405
LOAD MIF FILE INTO MIF DATABASE
410
USE MI TO OBTAIN MIF INFORMATION
415
420
PASS PREDETERMINED CRITERIA ?
NO
MAYBE
425
DOES USER WISH TO CONTINUE ?
NO
YES
YES
DOWNLOAD APPLET
430
435
PASS PREDETERMINED CRITERIA ?
MAYBE
440
DOES USER WISH TO CONTINUE ?
NO
YES
YES
INVOKE APPLET
445
END

FIG. 4

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 865 180 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
16.09.1998 Bulletin 1998/38

(51) Int Cl.[6]: **H04L 12/56**, H04L 29/06, H04L 29/12

(21) Application number: **98301577.7**

(22) Date of filing: **03.03.1998**

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **14.03.1997 US 818989**

(71) Applicant: **LUCENT TECHNOLOGIES INC.**
Murray Hill, New Jersey 07974-0636 (US)

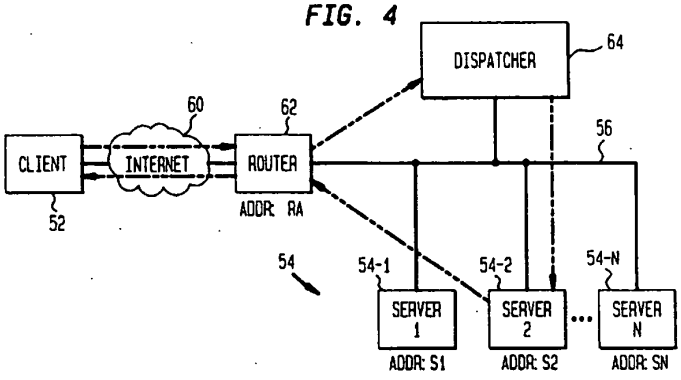
(72) Inventors:
- **Chung, Pi-Yu**
  Berkeley Heights, New Jersey 07922 (US)
- **Damani, Om P.**
  Austin, Texas 78705 (US)
- **Huang, Yennun**
  Bridgewater, New Jersey 08807 (US)
- **Kintala, Chandra M.**
  Warren, New Jersey 07059 (US)
- **Wang, Yi-Min**
  Berkeley Heights, New Jersey 07922 (US)

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr.
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(54) Load distribution among servers in a TCP/IP network

(57) Methods and apparatus for hosting a network service on a cluster of servers, each including a primary and a secondary Internet Protocol (IP) address. A common cluster address is assigned as the secondary address to each of the servers in the cluster. The cluster address may be assigned in UNIX-based servers using an *ifconfig alias* option, and may be a ghost IP address that is not used as a primary address by any server in the cluster. Client requests directed to the cluster address are dispatched such that only one of the servers of the cluster responds to a given client request. The dispatching may use a routing-based technique, in which all client requests directed to the cluster address are routed to a dispatcher connected to the local network of the server cluster. The dispatcher then applies a hash function to the client IP address in order to select one of the servers to process the request. The dispatching may alternatively use a broadcast-based technique, in which a router broadcasts client requests having the cluster address to all of the servers of the cluster over a local network. The servers then each provide a filtering routine, which may involve comparing a server identifier with a hash value generated from a client address, in order to ensure that only one server responds to each request broadcast by the router.

FIG. 4



EP 0 865 180 A2

Printed by Jouve, 75001 PARIS (FR)

Description

Field of the Invention

The present invention relates generally to data communication networks such as the Internet and more particularly to techniques for hosting network services on a cluster of servers used to deliver data over a network in response to client requests, where the cluster of servers can be collectively identified by a client using a single-address image.

Background of the Invention

With the explosive growth of the World Wide Web, many popular Internet web sites are heavily loaded with client requests. For example, it has been reported in S. L. Garfinkel, "The Wizard of Netscape," Webserver Magazine, July/August 1996, pp. 59-63, that home pages of Netscape Communications receive more than 80 million client requests or "hits" per day. A single server hosting a service is usually not sufficient to handle this type of aggressive growth. As a result, clients may experience slow response times and may be unable to access certain web sites. Upgrading the servers to more powerful machines may not always be cost-effective. Another common approach involves deploying a set of machines, also known as a cluster, and configuring the machines to work together to host a single service. Such a server cluster should preferably publicize only one server name for the entire cluster so that any configuration change inside the cluster does not affect client applications. The World Wide Web and other portions of the Internet utilize an application-level protocol, known as the Hypertext Transfer Protocol (HTTP), which is based on a client/server architecture. The HTTP protocol is described in greater detail in "Hypertext Transfer Protocol -- HTTP/1.0," Network Working Group, May 1996, <http://www.ics.uci.edu/pub/ietf/http>, which is incorporated by reference herein.

FIG. 1 illustrates an exemplary client/server architecture suitable for implementing HTTP-based network services on the Internet. A client 12 generates an HTTP request for a particular service, such as a request for information associated with a particular web site, and a Transmission Control Protocol/Internet Protocol (TCP/IP) connection is then established between the client 12 and a server 14 hosting the service. The client request is delivered to the server 14 in this example via a TCP/IP connection over a first network 16, a router 18 and a second network 20. The first network 16 may be a wide area communication network such as the Internet, while the second network 20 may be an Ethernet or other type of local area network (LAN) interconnecting server 14 with other servers in a server cluster. The router 18, also referred to as a gateway, performs a relaying function between the first and second networks which is transparent to the client 12.

The client request is generated by a web browser or other application-layer program operating in an application layer 22-1 of the client 12, and is responded to by a file transfer system or other program in an application layer 22-2 of the server 14. The requested network service may be designated by a Unifonn Resource Locator (URL) which includes a domain name identifying the server 14 or a corresponding server cluster hosting the service. The application-level program of the client 12 initiates the TCP/IP connection by requesting a local or remote Domain Name Service (DNS) to map the server domain name to an IP address. The TCP and IP packet routing functions in client 12 and server 14 are provided in respective TCP layers 24-1, 24-2 and IP layers 26-1, 26-2. The TCP and IP layers are generally associated with the transport and network layers, respectively, of the well-known Open Systems Interconnection (OSI) model. The TCP layers 24-1, 24-2 process TCP packets of the client request and server response. The TCP packets each include a TCP header identifying a port number of the TCP connection between the client 12 and server 14. The IP layers 26-1, 26-2 process IP packets formed from the TCP packets of the TCP layers. The IP packets each include an IP header identifying an IP address of the TCP/IP connection between the client 12 and server 14.

The IP address for a given network service may be determined, as noted above, by the client accessing a conventional DNS. The IP layer 26-1 of the client 12 uses the resulting IP address as a destination address in the IP packet headers of client request packets. The IP address together with the TCP port number provide the complete transport address for the HTTP server process. The client 12 and server 14 also include data link and physical layers 28-1 for performing framing and other operations to configure client request or reply packets for transmission over the networks 16 and 20. The router 18 includes data link and physical layers 28-3 for converting client request and server reply packets to IP format, and an IP layer 26-3 for performing packet routing based on IP addresses. The server 14 responds to a given client request by supplying the requested information over the established TCP/IP connection in a number of reply packets. The TCP/IP connection is then closed.

There are many known techniques for distributing HTTP client requests to a cluster of servers. FIGS. 2 and 3 illustrate server-side single-IP-address image approaches which present a single IP address to the clients. An example of this approach is the TCP router approach described in D.M. Dias, W. Kish, R. Mukherjee and R. Tewari, "A Scalable and Highly Available Web Server," Proceedings of COMPCON '96, pp. 85-92, 1996, which is incorporated by reference herein. FIG. 2 illustrates the TCP router approach in which a client 12 establishes a TCP/IP connection over Internet 30 with a server-side router 32 having an IP address RA. The router 32 is connected via a LAN 36 to a server cluster 34 including N servers 14-i, i = 1, 2, ... N, having respec-

tive IP addresses S1, S2, ... SN. Each server of the cluster 34 generally provides access to the same set of contents, and the contents may be replicated on a local disk of each server, shared on a network file system, or served by a distributed file system.

The single-address image is achieved by publicizing the address RA of the server-side router 32 to the clients via the DNS. The client 12 therefore uses RA as a destination IP address in its request. The request is directed to the router 32, which then dispatches the request to a selected server 14-k of server cluster 34 based on load characteristics, as indicated by the dashed line connecting client 12 to server 14-k via router 32. The router 32 performs this dispatching function by changing the destination IP address of each incoming IP packet of a given client request from the router address RA to the address Sk of selected server 14-k. The selected server 14-k responds to the client request by sending reply packets over the established TCP/IP connection, as indicated by the dashed line connecting server 14-k to client 12. In order to make the TCP/IP connection appear seamless to the client 12, the selected server 14-k changes the source IP address in its reply packets from its address Sk to the router address RA. The advantages of this approach are that it does not increase the number of TCP connections, and it is totally transparent to the clients. However, since the above-noted source IP address change is performed at the IP layer in a given server, the kernel code of every server in the cluster has to be modified to implement this mechanism A proposed hybrid of the DNS approach and the TCP router approach, in which a DNS server selects one of several clusters of servers using a round-robin technique, suffers from the same problem.

FIG. 3 illustrates a server-side single-address image approach known as network address translation, as described in greater detail in E. Anderson, D. Patterson and E. Brewer, "The Magicrouter, an Application of Fast Packet Interposing," Symposium on Operating Systems Design and Implementation, OSDI, 1996, <http://www.cs.berkeley.edu/~eanders/magicrouter/ osdi96-mr-submission.ps>, and Cisco Local Director, <http://www.cisco.com/warp/public/751/lodir/index.html>, which are incorporated by reference herein. As in the TCP router approach of FIG. 2, the client 12 uses the router address RA as a destination IP address in a client request, and the router 32 dispatches the request to a selected server 14-k by changing the destination IP address of each incoming request packet from the router address RA to the address Sk of selected server 14-k. However, in the network address translation approach, the source IP addresses in the reply packets from the selected server 14-k.are changed not by server 14-k as in FIG. 2, but are instead changed by the router 32. The reply packet flow indicated by a dashed line in FIG. 2 thus passes from server 14-k to client 12 via router 32.

Compared to the TCP router approach of FIG. 2, network address translation has the advantage of server transparency. That is, no specific changes to the kernel code of the servers are required to implement the technique. However, both the TCP router and network address translation approaches require that the destination address in a request packet header be changed to a server address so that the server can accept the request. These approaches also require that the source address in a reply packet header be changed to the router address so that the client can accept the reply. These changes introduce additional processing overhead and unduly complicate the packet delivery process. In addition, because of the address changes, the above-described single-address image approaches may not be suitable for use with protocols that utilize IP addresses within an application, such as that described in K. Egevang and P. Francis, "The IP Network Address Translator," Network Working Group, RFC 1631, <http://www.safety.net/ rfc1631.txt>, which is incorporated by reference herein. Furthermore, in both the TCP router and network address translation approaches, the router 32 needs to store an IP address mapping for every IP connection. Upon receiving an incoming packet associated with an existing TCP connection, the router has to search through all of the mappings to determine which server the packet should be forwarded to. The router itself may therefore become a bottleneck under heavy load conditions, necessitating the use of a more complex hardware design, as in the above-cited Cisco Local Director.

It is therefore apparent that a need exists for improved techniques for hosting a network service on a cluster of servers while presenting a single-address image to the clients, without the problems associated with the above-described conventional approaches.

Summary of the Invention

The present invention provides methods and apparatus for hosting a network service on a cluster of servers. All of the servers in a server cluster configured in accordance with the invention may be designated by a single cluster address which is assigned as a secondary address to each server. All client requests for a web site or other network service associated with the cluster address are sent to the server cluster, and a dispatching mechanism is used to ensure that each client request is processed by only one server in the cluster. The dispatching may be configured to operate without increasing the number of TCP/IP connections required for each client request. The invention evenly distributes the client request load among the various servers of the cluster, masks the failure of any server or servers of the cluster by distributing client requests to the remaining servers without bringing down the service, and permits additional servers to be added to the cluster without bringing down the service. Although well-suited for use in hosting web site services, the techniques of the present invention may also be used to support a wide variety of other

server applications.

In an exemplary embodiment of the invention, a network service is hosted by a server cluster in which each server includes a primary IP address and a secondary IP address. A common cluster address is assigned as the secondary IP address for each of the servers. The cluster address may be an IP address which does not correspond to a primary IP address of any of the servers. In UNIX-based servers, the cluster address may be assigned as the secondary address for a given server using an *ifconfig alias* option. If a given server includes multiple network interface cards, the cluster address may be assigned to one of the network interface cards using a UNIX *ifconfig* command without the *alias* option, or other similar technique. A router is coupled to a local network of the server cluster and is also coupled via the Internet to a client. The router receives client requests from the Internet, and uses a dispatching technique to direct client requests having the cluster address as a destination. The client requests are dispatched such that each of the requests is processed by only one of the servers in the cluster. The dispatching function may be based on the result of applying a hash function to an IP address of the given client. A suitable hash function may be determined using an analysis of a distribution of client IP addresses in an access log associated with one or more of the servers. In the event that a server has failed, the hash function may be reapplied to the client IP address to identify another server.

Two illustrative dispatching techniques for providing a single-address image for a server cluster in accordance with the invention include routing-based dispatching and broadcast-based dispatching. In the routing-based technique, a dispatcher is coupled to the router and to a local network of the server cluster. The router directs client requests having the cluster address to the dispatcher, and the dispatcher selects a particular one of the servers to process a given client request based on the result of applying a hash function to the client address. In the broadcast-based technique, the router broadcasts client requests having the cluster address to each of the servers over the local network of the server cluster. Each of the servers implements a filtering routine to ensure that each client request is processed by only one of the servers. The filtering routine may involve applying a hash function to the client IP address associated with a given client request, and comparing the result to a server identifier to determine whether that server should process the client request.

The techniques of the present invention provide fast dispatching and can be implemented with reduced cost and complexity. The techniques are suitable for use in TCP/IP networks as well as networks based on a variety of other standards and protocols. Unlike the conventional single-address image approaches, the present invention does not require that a destination address in a request packet header be changed to a server address so that the server can accept the request, or that a source address in a reply packet header be changed to the router address so that the client can accept the reply. In addition, the router need not store an IP address mapping for every IP connection, nor is it required to search through such a mapping to determine which server a packet should be forwarded to. The router itself will therefore not become a bottleneck under heavy load conditions, and special router hardware designs are not required. These and other features and advantages of the present invention will become more apparent from the accompanying drawings and the following detailed description.

Brief Description of the Drawings

FIG. 1 is a block diagram illustrating a conventional client-server interconnection in accordance with the TCP/IP standard;

FIG. 2 illustrates a prior art TCP router technique for hosting a network service on a cluster of servers;

FIG. 3 illustrates a prior art network address translation technique for hosting a network service on a cluster of servers;

FIG. 4 illustrates a technique for hosting a network service on a cluster of servers using routing-based dispatching in accordance with an exemplary embodiment of the invention; and

FIG. 5 illustrates a technique for hosting a network service on a cluster of servers using broadcast-based dispatching in accordance with another exemplary embodiment of the invention.

Detailed Description of the Invention

The present invention will be illustrated below in conjunction with exemplary client/ server connections established over the Internet to a server cluster using the Transmission Control Protocol/Internet Protocol (TCP/IP) standard. It should be understood, however, that the invention is not limited to use with any particular type of network or network communication protocol. The disclosed techniques are suitable for use with a wide variety of other networks and protocols. The term "server cluster" as used herein refers to a group or set of servers interconnected or otherwise configured to host a network service. The terms "cluster address" and "single-address image" refer generally to an address associated with a group of servers configured to support a network service or services. A "ghost IP address" is one type of cluster address in the form of an IP address which is not used as a primary address for any server of a given server cluster. The term "network service" is intended to include web sites, Internet sites and data delivery services, as well as any other data transfer -mechanism accessible by a client over a network. The term "client request" refers to a communication from a client which initiates the network service. A given client request may include multiple packets or only a single

packet, depending on the nature of the request.

The present invention provides an improved single-address image approach to distributing client requests to servers of a server cluster. In a preferred embodiment, the invention allows all servers of a server cluster to share a single common IP address as a secondary address. The secondary address is also referred to herein as a cluster address, and may be established using an *ifconfig alias* option available on most UNIX-based systems, or similar techniques available on other systems. The cluster address may be publicized to clients using the above-noted Domain Name Service (DNS) which translates domain names associated with Uniform Resource Locators (URLs) to IP addresses. All client requests to be directed to a service hosted by the server cluster are sent to the single cluster address, and dispatched to a selected one of the servers using routing-based or broadcast-based dispatching techniques to be described in greater detail below. Once a server is selected, future request packets associated with the same client request may be directed to the same server. All other communications within the server cluster may utilize primary IP addresses of the servers.

The above-noted *ifconfig alias* option is typically used to allow a single server to serve more than one domain name. For example, the *ifconfig alias* option allows a single server to attach multiple IP addresses, and thus multiple domain names, to a single network interface, as described in "Two Servers, One Interface" <http://www.thesphere.com/~dlp/TwoServers/>, which is incorporated by reference herein. Client requests directed to any of the multiple domain names can then be serviced by the same server. The server determines which domain name a given request is associated with by examining the destination address in the request packet. The present invention utilizes the *ifconfig alias* option to allow two servers to share the same IP address. Normally, two servers cannot share the same IP address because such an arrangement would cause any packet destined for the shared address to be accepted and responded to by both servers, confusing the client and possibly leading to a connection reset. Therefore, before a server is permitted to attach a new IP address to its network interface, a check may be made to ensure that no other server on the same local area network (LAN) is using that IP address. If a duplicate address is found, both servers are informed and warnings are issued. The routing-based or broadcast-based dispatching of the present invention ensures that every packet is processed by only one server of the cluster, such that the above-noted warnings do not create a problem.

An alternative technique for assigning a secondary address to a given server of a server cluster in accordance with the invention involves configuring the given server to include multiple network interface cards such that a different address can be assigned to each of the network interface cards. For example, in a UNIX-based system, conventional *ifconfig* commands may be used, without the above-described *alias* option, to assign a primary IP address to one of the network interface cards and a secondary IP address to another of the network interface cards. The secondary IP address is also assigned as a secondary IP address to the remaining servers in the cluster, and used as a cluster address for directing client requests to the cluster.

The exemplary embodiments of the present invention to be described below utilize dispatching techniques in which servers are selected based on a hash value of the client IP address. The hash value may be generated by applying a hash function to the client IP address, or by applying another suitable function to generate a hash value from the client IP address. For example, given N servers and a packet from a client having a client address CA, a dispatching function in accordance with the invention may compute a hash value *k* as CA *mod (N-1)* and select server *k* to process the packet. This ensures that all request or reply packets of the same TCP/IP connection are directed to the same server in the server cluster. A suitable hash function may be determined by analyzing a distribution of client IP addresses in actual access logs associated with the servers such that client requests are approximately evenly distributed to all servers. When a server in the cluster fails, the subset of clients assigned to that server will not be able to connect to it. The present invention addresses this potential problem by dynamically modifying the dispatching function upon detection of a server failure. If the hash value of a given client IP address maps to the failed server, the client IP address is rehashed to map to a non-failed server, and the connections of the remaining clients are not affected by the failure.

FIG. 4 illustrates a routing-based dispatching technique in accordance with the present invention. Solid lines indicate network connections, while dashed lines show the path of an exemplary client request and the corresponding reply. A client 52 sends a client request to a server cluster 54 including N servers 54-i, i = 1, 2, ... N having IP addresses S1, S2, ... SN and interconnected by an Ethernet or other type of LAN 56. The client request is formulated in accordance with the above-described HTTP protocol, and may include a URL with a domain name associated with a web site or other network service hosted by the server cluster 54. The client accesses a DNS to determine an IP address for the domain name of the service, and then uses the IP address to establish a TCP/IP connection for communicating with one of the servers 54-i of the server cluster 54. In accordance with the invention, a "ghost" IP address is publicized to the DNS as a cluster address for the server cluster 54. The ghost IP address is selected such that none of the servers 54-i of cluster 54 has that IP address as its primary address. Therefore, any request packets directed to the ghost IP address are associated with client requests for the service of the single-address image cluster 54. The use of the ghost IP address thus distin-

guishes a network service hosted by the cluster from activities of the servers 54-i which utilize the primary server addresses, and prevents interference with these primary address activities.

The client 52 uses the ghost IP address as a cluster address for directing its request to the server cluster 54. The request is directed over Internet 60 to a router 62 having an IP address RA. The router 62 includes a routing table having an entry or record directing any incoming request packets having the ghost IP address to a dispatcher 64 connected to the LAN 56. The dispatcher 64 includes an operating system configured to run in a router mode, using a routing algorithm which performs the dispatching described herein. In alternative embodiments, the functions of the dispatcher 64 could be incorporated into the router 62 in order to provide additional efficiency improvements. Each of the servers 54-i of the cluster 54 utilizes the above-described *ifconfig alias* option to set the ghost IP address as their secondary address. As noted above, this technique for setting a secondary address for each of the servers 54-i generally does not require any alteration of the kernel code running on the servers. In alternative embodiments, one or more of the servers 54-i may be configured to include multiple network interface cards, as previously noted, such that a different address can be assigned to each of the network interface cards of a given server using a UNIX *ifconfig* command or other similar technique.

The router 62 routes any packets having the ghost IP address to the dispatcher 64 in accordance with the above-noted routing table record. The dispatcher 64 then applies a hash function to the client IP address in a given request packet to determine which of the servers 54-i the given packet should be routed to. In the example illustrated in FIG. 4, the dispatcher 64 applies a hash function to the IP address of client 52 and determines that the corresponding request packet should be routed to server 54-2 at IP address S2. The dispatcher 64 then routes the request packet to the server 54-2 over LAN 56, as indicated by the dashed line, using the primary address S2 of server 54-2 to distinguish it from the other servers of cluster 54. After the network interface of server 54-2 accepts the packet, all higher level processing may be based on the ghost IP address because that is the destination address in the packet IP header and possibly in the application-layer packet contents. After processing the request, the server 54-2 replies directly to the client 52 via router 62 over the established TCP/IP connection, using the ghost IP address, and without passing through the dispatcher 64.

It should be noted that when a request packet destined for the ghost IP address is received by the network interface of dispatcher 64 and placed back onto the same network interface for delivery to one of the servers 54-i over LAN 56, it may cause an Internet control message protocol (ICMP) host redirect message to be sent to the router 62. This ICMP message is designed to direct the router 62 to update its routing table such that any future packets having the ghost IP address can bypass the dispatcher 64 and go directly to the destination server, as described in greater detail in W.R. Stevens, TCP/IP Illustrated, Vol. 1, Ch. 6, pp. 69-83, which is incorporated by reference herein. However, this effect is undesirable in the routing technique of FIG. 4 because the dispatcher 64 performs the server selection process as previously described. It therefore may be necessary to suppress the ICMP host redirect message for the ghost IP address by, for example, removing or altering the corresponding operating system code in the dispatcher. In the above-mentioned alternative embodiments in which the dispatching function is implemented within the router 62, the ICMP redirect message is not generated and therefore need not be suppressed. Another potential problem may arise when a reply packet is sent back to the client 52 from the selected server 54-2 with the ghost IP address, in that it may cause the router 62 to associate, in its Address Resolution Protocol (ARP) cache, the ghost IP address with the LAN address of the selected server. The operation of the ARP cache is described in greater detail in W.R. Stevens, TCP/IP Illustrated, Vol. 1, Chs. 4 and 5, pp. 53-68, which is incorporated by reference herein. The illustrative embodiment of FIG. 4 avoids this problem by automatically routing the request packets to the dispatcher 64, and then dispatching based on the server primary IP address, such that the router ARP cache is not used.

FIG. 5 illustrates a broadcast-based dispatching technique in accordance with the present invention. Again, solid lines indicate network connections, while dashed lines show the path of an exemplary client request and the corresponding reply. As in the FIG. 4 routing-based embodiment, client 52 sends a client request to server cluster 54 including N servers 54-i, i = 1, 2, ... N connected to LAN 56 and having IP addresses S1, S2, ... SN. The client 52 uses the above-described ghost address as a cluster address for directing its request to the server cluster 54. The request is directed over Internet 60 to a router 70 having an IP address RA. The router 70 broadcasts any incoming request packets having the ghost IP address to the LAN 56 interconnecting the servers 54-i of the server cluster 54, such that the request packet is received by each of the servers 54-i.

Each of the servers 54-i of the cluster 54 implements a filtering routine in order to ensure that only one of the servers 54-i processes a given client request. The filtering routine may be added to a device driver of each of the servers 54-i. In an exemplary implementation, each of the servers 54-i is assigned a unique identification (ID) number. The filtering routine of a given server 54-i computes a hash value of the client IP address and compares it to the ID number of the given server. If the hash value and the ID number do not match, the filtering routine of the given server rejects the packet. If the hash value and the ID number do match, the given server accepts and processes the packet as if it had received the packet through a conventional IP routing mechanism. In

the illustrative example of FIG. 5, a packet associated with request from client 52 is broadcast by the router 70 to each of the servers 54-i of the server cluster 54 over the LAN 56 as previously noted. The filtering routine of server 54-2 generates a hash value of the client IP address which matches the unique ID number associated with server 14-2, and server 14-2 therefore accepts and processes the packet. The filtering routines of the N-1 other servers 54-i each indicate no match between the client IP address and the corresponding server ID number, and therefore discard the broadcast packet. The reply packets are sent back to the client 52 via router 70, as indicated by the dashed lines, using the ghost IP address.

The broadcast-based dispatching technique of FIG. 5 may be implemented using a permanent ARP entry within the router 70, to associate the ghost IP address with the Ethernet or other local network broadcast address associated with LAN 56 of the cluster 54. A potential problem is that any reply packet from a selected server appears to be coming from the ghost IP address, and may therefore cause the router 70 to overwrite the entry in its ARP cache such that the ghost IP address is associated with the LAN address of the selected server. This potential problem may be addressed by setting up a routing table entry in the router 70 to direct all packets having a ghost IP destination address to a second ghost IP address which is a legal subnet address in the LAN 56 of the server cluster 54 but is not used by any server. In addition, an entry is inserted in the ARP cache of the router 70 to associate the second ghost IP address with the broadcast address of the LAN 56. When the router 70 routes a packet to the second ghost IP address, it will then actually broadcast the packet to each of the servers 54-i of the cluster 54. Since no reply packet is sent from the second ghost IP address, the corresponding entry of the router ARP cache will remain unchanged. Another potential problem is that some operating systems, such as the NetBSD operating system, do not allow a TCP packet to be processed if it is received from a broadcast address. This potential problem may be avoided by a suitable modification to the broadcast address in the LAN packet header attached to the packet.

The routing-based and broadcast-based dispatching techniques described in conjunction with FIGS. 4 and 5 above have been implemented on a cluster of Sun SPARC workstations. The NetBSD operating system, as described in NetBSD Project, <http:/www. NetBSD. org>, was used to provide any needed kernel code modifications. The dispatching overhead associated with both techniques is minimal because the packet dispatching is based on simple IP address hashing, without the need for storing or searching any address-mapping information. In the routing-based dispatching technique, the additional routing step in the dispatcher 64 typically adds a delay of about 1 to 2 msecs to the TCP round-trip time of each incoming request packet. A study in W. R. Stevens, TCP/IP Illustrated, Volume 3, pp. 185-186, which is incorporated by reference herein, indicates that the median TCP round-trip time is 187 msecs. The additional delay attributable to the routing-based dispatching is therefore negligible. Although the additional routing step for every request packet sent to the ghost IP address may increase the traffic in the LAN of the server cluster, the size of a request in many important applications is typically much smaller than that of the corresponding response, which is delivered directly to the client without the additional routing. In the broadcast-based dispatching technique, the broadcasting of each incoming request packet on the LAN of the server cluster does not substantially increase network traffic. Although a hash value is computed for each incoming packet having the ghost IP destination address, which increases the CPU load of each server, this additional computation overhead is negligible relative to the corresponding communication delay.

Both the routing-based and broadcast-based dispatching techniques of the present invention are scalable to support relatively large numbers of servers. Although the dispatcher in the routing-based technique could present a potential bottleneck in certain applications, a study in the above-cited D.M. Dias et al. reference indicates that a single dispatcher can support up to 75 server nodes, which is sufficient support for many practical systems. The number of servers supported may be even higher with the present invention given that the routing-based dispatching functions described herein are generally simpler than those in the D.M. Dias et al. reference. It should also be noted that additional scalability can be obtained by combining the routing-based dispatching of the present invention with a DNS round-robin technique. For example, a DNS server may be used to map a domain name to one of a number of different ghost IP addresses belonging to different server clusters using a round-robin technique. In the broadcast-based dispatching technique, there is no potential dispatching bottleneck, although the device drivers or other portions of the servers may need to be modified to provide the above-described filtering routines.

The routing-based and broadcast-based dispatching of the present invention can also provide load balancing and failure handling capabilities. For example, given N servers and a packet from client address CA, the above-described routing-based dispatching function may compute a hash value k as $CA \bmod (N-1)$ and select server k to process the packet. More sophisticated dispatching functions can also be used, and may involve analyzing the actual service access log to provide more effective load balancing. In order to detect failures, each server may be monitored by a watchdog daemon such as the *watchd* daemon described in greater detail in Y. Huang and C. Kintala, "Software Implemented Fault Tolerance: Technologies and Experience," Proceedings of the 23rd International Symposium on Fault-Tolerant Computing - FTCS, Toulouse, France, pp. 2-9, June

1993, which is incorporated by reference herein. When a server fails, the corresponding *watchd* daemon initiates a change of the dispatching function to mask the failure and rebalance the load. A system call interface may be implemented to allow the dispatching function to be changed while the servers remain on-line. In routing-based dispatching, the *watchd* daemon may notify the dispatcher to change the dispatching function, while in broadcast-based dispatching, all servers may be notified to modify their filtering routines. For example, if a server *k* fails, the new dispatching function may check to see if the hash value *CA mod N* equals *k*. If it does, a new hash value *j = CA mod (N-1)* is computed. If *j* is less than *k*, the packet goes to server *j*. Otherwise, the packet goes to server *j+1*. This technique does not affect the clients of non-failed servers, reassigns the clients of the failed server evenly to the remaining servers, and can be readily extended to handle multiple server failures. Additional servers can be added to the cluster without bringing down the service by changing the dispatching function from *CA mod N* to *CA mod (N+1)*.

In routing-based dispatching, the dispatcher may become a single point of failure, and therefore should also be monitored by a *watchd* daemon or other suitable failure monitoring mechanism. Upon detecting a failure, the *watchd* daemon may trigger a transfer of the dispatching function from the primary dispatcher to a backup dispatcher, and then direct the router to change the entry in its routing table such that future incoming request packets are routed to the backup dispatcher. Since no mapping table is maintained by the primary dispatcher, this approach is substantially stateless. Proper routing may be ensured by simply utilizing consistent routing functions in the primary and backup dispatchers, without the substantial additional costs associated with mapping-based approaches.

The use of the *ifconfig alias* option or other similar technique to provide a single-address image for a server cluster provides a number of advantages over the conventional techniques described previously. For example, it avoids the need to change the destination address in a request packet header so that a particular server can accept the request, and the need to change the source address in a reply packet header to the cluster address so that the client can accept the reply. With the single-address image approach of the present invention, all servers can accept and respond to packets having the cluster address, so that the addresses in the request and reply packet headers do not need to be modified. Since the single-image approach of the present invention does not require alternation of the packet addresses, it is suitable for use with a wide variety of protocols, including those protocols which utilize IP addresses within an application program. In addition, the single-address image approach of the present invention does not require a router to store or to search through a potentially large number of IP address mappings in order to determine which cluster server should receive a request packet. The invention thus effectively removes the possibility that the router may become a bottleneck under heavy load conditions.

The above-described embodiments of the invention are intended to be illustrative only. Numerous alternative embodiments may be devised by those skilled in the art without departing from the scope of the following claims.

Claims

1. A method of routing client requests to a plurality of servers configured to support a network service over a communication network, each of the servers having a primary address, the method comprising the steps of:

   assigning a common address as a secondary address for each of the plurality of servers; and processing client requests directed to the common address such that each of the requests is processed by a particular one of the plurality of servers.

2. The method of claim 1 wherein the network utilizes a TCP/IP protocol and the primary and secondary addresses are primary and secondary IP addresses, respectively.

3. The method of claim 2 wherein the common address is an IP address which does not correspond to a primary IP address of any of the plurality of servers.

4. The method of any of the preceding claims wherein at least one of the plurality of servers is a UNIX-based server including multiple network interface cards, and the assigning step includes assigning the common address for the at least one server using an *ifconfig* command..

5. The method of any of claims 1 to 3 wherein the plurality of servers are UNIX-based servers, and the assigning step includes assigning the common address utilizing an *ifconfig alias* option for at least a subset of the plurality of servers.

6. The method of any of the preceding claims wherein the processing step includes the step of dispatching a request of a given client to one of the plurality of servers based on application of a hash function to an IP address of the given client.

7. The method of claim 6 wherein the hash function is determined based on an analysis of a distribution of client IP addresses in an access log associated with one or more of the servers.

8. The method of claim 6 wherein the dispatching step includes reapplying the hash function to the client IP address to identify another servers if a server identified as a results of a previous application of the hash function has failed.

9. The method of any of claims 1 to 5 wherein the processing step includes the steps of:

routing client requests directed to the common address to a dispatcher connected to a local network associated with the plurality of servers; and
selecting a particular one of the servers to process a given client request based on application of a hash function to a corresponding client address in the dispatcher.

10. The method of any of claims 1 to 5 wherein the processing step includes the steps of:

broadcasting a given client request directed to the common address to each of the plurality of servers over a local network associated with the servers; and
implementing a filtering routine in each of the plurality of servers so that the given client request is processed by only one of the servers.

11. The method of claim 10 wherein the implementing step includes the steps of:

applying a hash function to a client IP address associated with the given client request; and
comparing the result of the applying step to an identifier of a particular server to determine whether that server should process the given client request.

12. An apparatus for routing client requests to a plurality of servers configured to support a network service over a communication network, each of the servers having a primary address, the apparatus comprising:

means for assigning a common address as a secondary address for each of the plurality of servers; and
means for processing client requests directed to the common address such that each of the requests is processed by a particular one of the plurality of servers.

13. The apparatus of claim 12 wherein the processing means is operative to dispatch a request of a given client to one of the plurality of servers based on application of a hash function to an IP address of the given client.

14. The apparatus of claim 13 wherein the hash function is determined based on an analysis of a distribution of client IP addresses in an access log associated with one or more of the servers.

15. The apparatus of claim 13 wherein the processing means is further operative to reapply the hash function to the client IP address to identify another server if a server identified as a result of a previous application of the hash function has failed.

16. The apparatus of any of claims 12 to 15 wherein the processing means further includes a dispatcher connected to a local network associated with the plurality of servers, wherein the dispatcher is operative to receive client requests directed to the common address, and to select a particular one of the servers to process a given client request based on application of a hash function to a corresponding client address.

17. The apparatus of any of claims 12 to 15 wherein the processing means further includes:

means for broadcasting a given client request directed to the common address to each of the plurality of servers over a local network associated with the servers; and
means for filtering the given client request in each of the plurality of servers so that the given client request is processed by only one of the servers.

18. The apparatus of claim 17 wherein the filtering means is operative to apply a hash function to a client IP address associated with the given client request, and to compare the result of the applying step to an identifier of a particular server to determine whether that server should process the given client request.

19. An apparatus for routing client requests for a network service over a communication network, the apparatus comprising:

a plurality of servers configured to support the network service, each of the servers having a primary address and a secondary address, wherein a common address is assigned as the secondary address for each of the plurality of servers; and
a router coupled to the servers and operative to route client requests directed to the common address such that each of the requests is processed by a particular one of the plurality of servers.

20. The apparatus of claim 19 wherein the router is fur-

ther operative to route client requests such that a request of a given client is routed to one of the plurality of servers based on application of a hash function to an IP address of the given client.

21. The apparatus of claim 20 wherein the hash function is determined based on an analysis of a distribution of client IP addresses in an access log associated with one or more of the servers.

22. The apparatus of claim 20 wherein the hash function is reapplied to the client IP address to identify another server if a server identified as a result of a previous application of the hash function has failed.

23. The apparatus of any of claims 19 to 22 further including a dispatcher coupled to the router and to a local network associated with the plurality of servers, such that the router directs client requests having the common address to the dispatcher, and the dispatcher selects a particular one of the servers to process a given client request based on application of a hash function to a corresponding client address.

24. The apparatus of any of claims 19 to 22 wherein the router is further operative to broadcast a given client request directed to the common address to each of the plurality of servers over a local network associated with the servers, and further wherein each of the servers implements a filtering routine so that the given client request is processed by only one of the servers.

25. The apparatus of claim 24 wherein the filtering routine involves applying a hash function to a client IP address associated with the given client request. and comparing the result to an identifier of a particular server to determine whether that server should process the given client request.

26. The apparatus of any of claims 12 to 25 wherein the network utilizes a TCP/IP protocol and the primary and secondary addresses are primary and secondary IP addresses, respectively.

27. The apparatus of claim 26 wherein the common address is an IP address which does not correspond to a primary IP address of any of the plurality of servers.

28. The apparatus of any of claims 12 to 27 wherein at least one of the plurality of servers is a UNIX-based server including multiple network interface cards, and the common address is assigned for the at least one server using an *ifconfig* command.

29. The apparatus of any of claims 12 to 27 wherein the plurality of servers are UNIX-based servers, and the common address is assigned as the secondary address of the plurality of servers by utilizing an *ifconfig alias* option for at least a subset of the plurality of servers.

FIG. 1
(PRIOR ART)

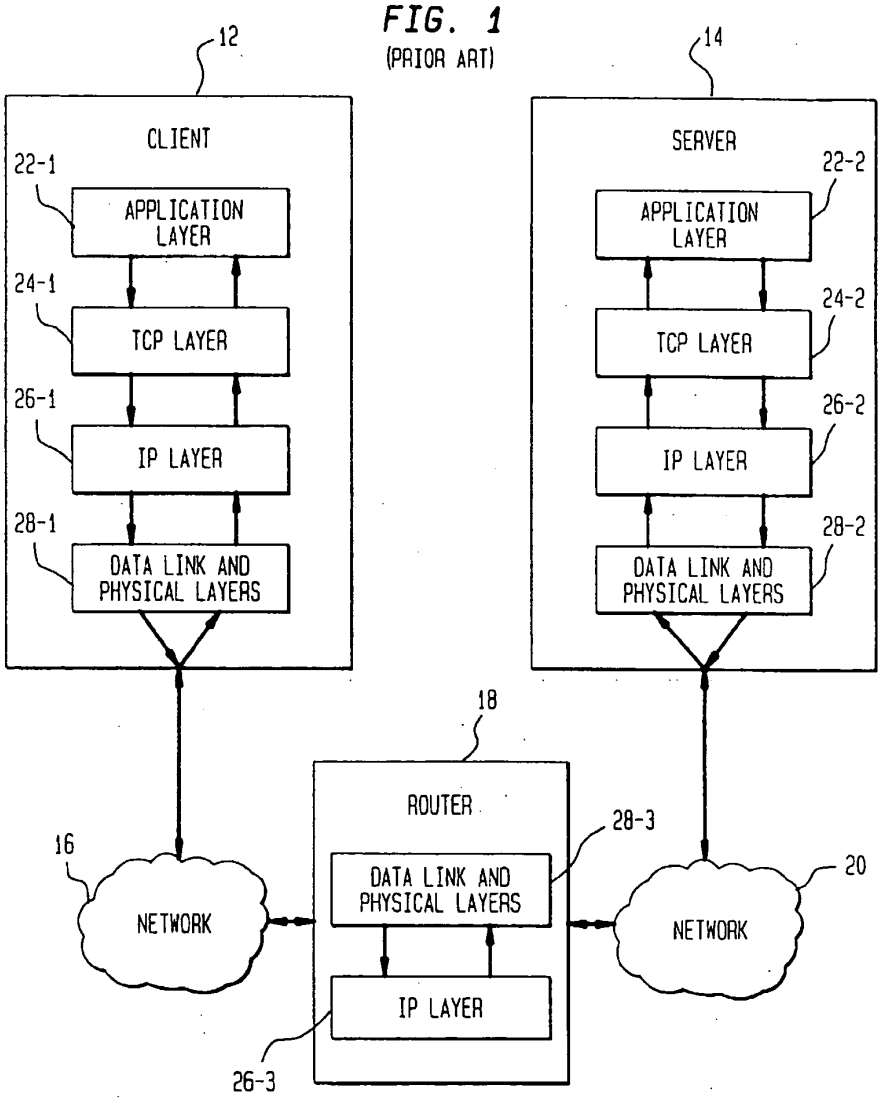


FIG. 2
(PRIOR ART)

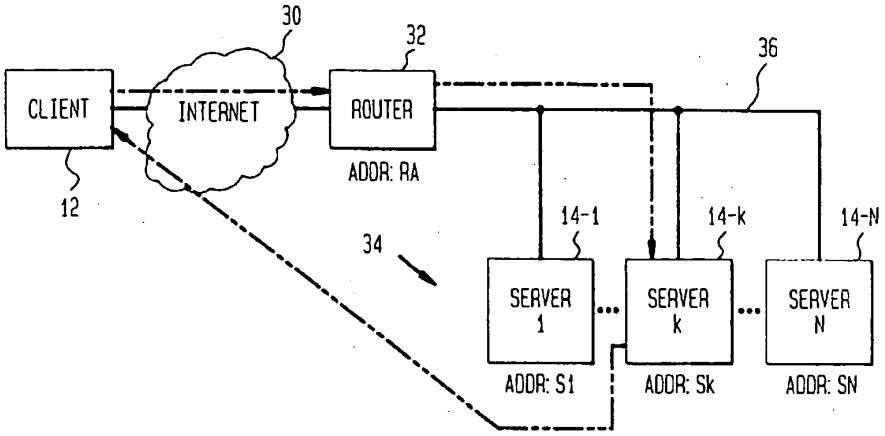


FIG. 3
(PRIOR ART)

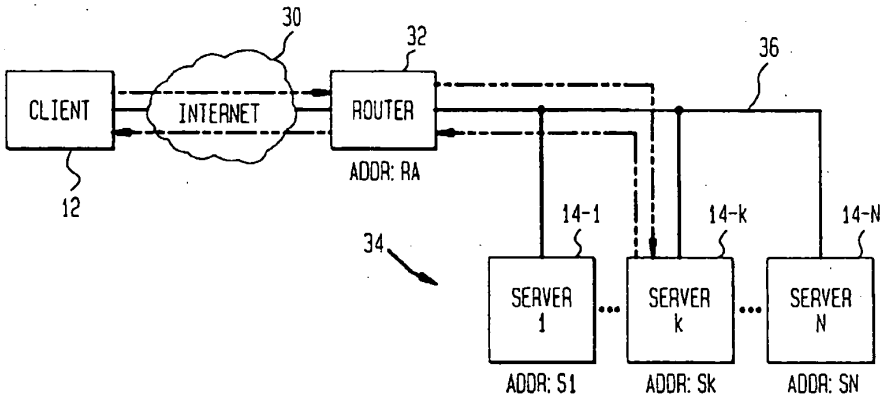


FIG. 4

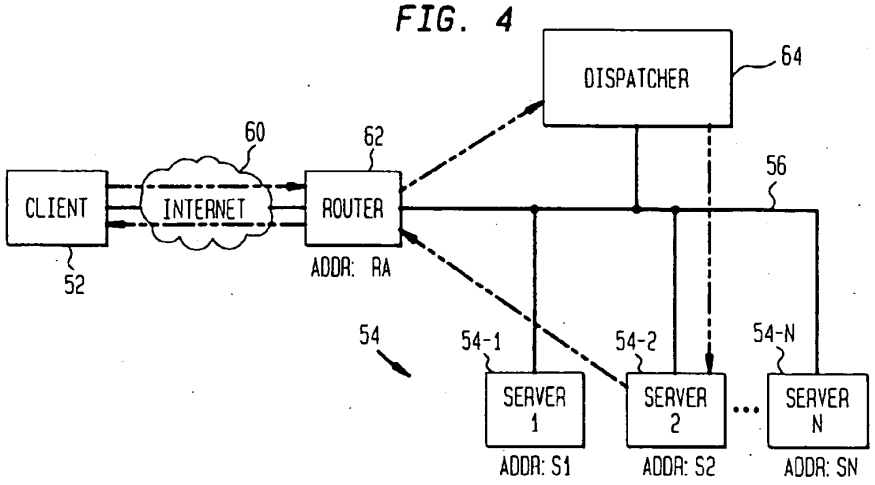
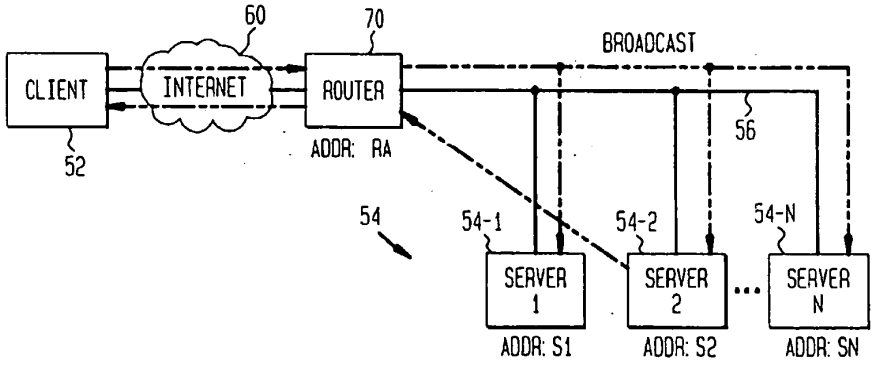

DISPATCHER
64
60
62
CLIENT
INTERNET
ROUTER
56
52
ADDR: RA
54
54-1
54-2
54-N
SERVER 1
SERVER 2
SERVER N
ADDR: S1
ADDR: S2
ADDR: SN

FIG. 5


60
70
BROADCAST
CLIENT
INTERNET
ROUTER
56
52
ADDR: RA
54
54-1
54-2
54-N
SERVER 1
SERVER 2
SERVER N
ADDR: S1
ADDR: S2
ADDR: SN

OIPE IAP97
JAN 25 2008
PATENT & TRADEMARK OFFICE

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2142
Application Serial No.: 11/065,412	Examiner: BLAIR, Douglas B.
Application Filing Date: February 23, 2005	Confirmation No.: 3296
Title: **DELIVERING RESOURCES IN A SHARED CONTENT DELIVERY INFRASTRUCTURE**	Date: January 25, 2008

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached (except for copies not required to be submitted).

It is respectfully requested that the cited references be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom. It is further respectfully requested that the enclosed form(s) PTO-1449 be initialed by the Examiner to indicate such consideration, and that a copy thereof be returned to applicant.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action

which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed within three (3) months of the U.S. filing date OR before the mailing date of a first Office Action on the merits. No certification or fee is required.

RELATED APPLICATIONS

The Examiner's attention is directed to the following related applications:

09/021,506 (filed 02-10-1998, now U.S. Patent No. 6,185,598);
09/612,598 (filed 07-07-2000, abandoned);
09/930,975 (filed 08-17-2001, abandoned);
10/004,430 (filed 11-25-2003, now U.S. Patent No. 6,554,807);
10/095,811 (filed 03-13-2002, now U.S. Patent No. 7,054,935);
11/441,253 (filed 05-26-2006, Pending);
11/806,147 (filed 05-30-2007, Pending);
11/806,152 (filed 05-30-2007, Pending);
11/806,153 (filed 05-30-2007, Pending);
11/806,154 (filed 05-30-2007, Pending);
11/980,686 (filed 10-31-2007, Pending);
11/980,356 (filed 10-31-2007, Pending);
11/980,689 (filed 10-31-2007, Pending);
11/980,684 (filed 10-31-2007, Pending);and
11/980,401 (filed 10-31-2007, Pending).

All of the above applications are directed to related technical subject matter. The identification of these U.S. Patent Applications is not to be construed as a waiver of secrecy as to those applications now or upon issuance of the present application as a patent. The Examiner is respectfully requested to consider art cited therein during examination.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 42624	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430



INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified)	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2142
	Examiner Name	BLAIR, Douglas B.
	Attorney Docket No.	2711-0073
Sheet 1 of 1	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
	1-1	Complaint for Injunctive and Other Relief, filed 12/17/2007 in Civil Action No. 2:07CV589	
	1-2	Report on the Filing or Determination of an Action Regarding a Patent or Trademark, dated 12/28/2007, in connection with Civil Action No. 2:07CV589 filed 12/17/2007.	
	1-3	Rule 7.1 Financial Interest Disclosure Statement filed 12/17/2007 in Civil Action No. 2:07CV589	
	1-4		
	1-5		
	1-6		
	1-7		
	1-8		
	1-9		
	1-10		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:	Atty. Docket No.: 2711-0073
FARBER, David et al.	Confirmation No.: 3296
Appln. No.: 11/065,412	Group Art Unit: 2142
Filed: February 23, 2005	Examiner: Blair, Douglas
Title: SHARED CONTENT DELIVERY INFRASTRUCTURE (as amended)	Date: January 17, 2008

LETTER RE LITIGATION IN RELATED PATENTS

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This patent application is related to U.S. Patent Nos. 7,054,935 and 6,654,807, both of which have been recently asserted in a litigation for patent infringement in the Eastern District of Virginia, in a case captioned: *Level 3 Communications, LLC v. Limelight Networks, Inc.* (2:07-cv-00589-WDK-FBS). The case was filed on 12/17/2007.

U.S. Patent No. 6,473,405 was also asserted in that litigation.

The Examiner is kindly requested to contact the undersigned with any questions about this litigation.

CUSTOMER NUMBER
42624

Respectfully submitted,

By: /Brian Siritzky/Reg.No. 37497
Brian Siritzky, Ph.D.
Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor,
Arlington Virginia 22203
Main: (703) 894-6400 ● FAX: (703) 894-6430

Electronic Acknowledgement Receipt

EFS ID:	2733212
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	SHARED CONTENT DELIVERY INFRASTRUCTURE
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	42624
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	17-JAN-2008
Filing Date:	23-FEB-2005
Time Stamp:	16:47:33
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes) /Message Digest	Multi Part /.zip	Pages (if appl.)
1	Miscellaneous Incoming Letter	PA_Letter_Re_Litigation_In_Related_Patents.pdf	117877 a624782dabdbb657f178931c23ff8b7ecdd24c71	no	1

Warnings:	
Information:	

Total Files Size (in bytes):	117877

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

OIPE 1AP44
JUL 12 2007
PATENT & TRADEMARK OFFICE

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:	Attorney Docket: 2711-0073
FARBER, David et al.	Group Art Unit: 2157
Application Serial No.: 11/065,412	Examiner: Unassigned
Application Filing Date: February 23, 2005	Date: July 12, 2007
Title: **OPTIMIZED NETWORK RESOURCE LOCATION**	Confirmation No.: 3296

INFORMATION DISCLOSURE STATEMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to 37 C.F.R. § 1.56, the attention of the Patent and Trademark Office is hereby directed to the reference(s) listed on the attached PTO-1449. One copy of each non-U.S. Patent reference is attached (except for copies not required to be submitted).

Copies of some the references were submitted to, or cited by, the Office in parent Application No. 09/612,598, filed July 7, 2000, and/or grandparent Application No. 09/021,506, filed February 10, 1998, which are relied upon for an earlier effective filing date under 35 U.S.C. 120, and copies of the references were furnished at that time. Accordingly, copies of those references are not submitted herewith. 37 C.F.R. §1.98(d).

It is respectfully requested that the cited references be expressly considered during the prosecution of this application, and that the reference(s) be made of record therein and appear among the "References Cited" on any patent to issue therefrom. It is further respectfully requested that the enclosed form(s) PTO-1449

be initialed by the Examiner to indicate such consideration, and that a copy thereof be returned to applicant.

The submission of any document herewith, which is not a statutory bar, is not intended that any such document constitutes prior art against any of the claims of the present application or is considered to be material to patentability as defined in 37 C.F.R. § 1.56(b). Applicants do not waive any rights to take any action which would be appropriate to antedate or otherwise remove as a competent reference against the claims of the present application.

This Information Disclosure Statement is being filed within three (3) months of the U.S. filing date OR before the mailing date of a first Office Action on the merits. No certification or fee is required.

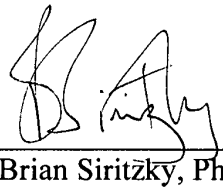
RELATED APPLICATIONS

The Examiner's attention is directed to the following related applications:

09/021,506 (filed 02-10-1998, now U.S. Patent No. 6,185,598);
09/612,598 (filed 07-07-2000, abandoned);
09/930,955 (filed 08-17-2001, abandoned);
10/004,430 (filed 11-25-2003, now U.S. Patent No. 6,554,807);
10/095,811 (filed 03-13-2002, now U.S. Patent No. 7,054,935);
11/441,253 (filed 05-26-2006, Pending);
11/806,147 (filed 05-30-2007, Pending);
11/806,152 (filed 05-30-2007, Pending);
11/806,153 (filed 05-30-2007, Pending); and
11/806,154 (filed 05-30-2007, Pending).

All of the above applications are directed to related technical subject matter. The identification of these U.S. Patent Applications is not to be construed as a waiver of secrecy as to those applications now or upon issuance of the present application as a patent. The Examiner is respectfully requested to consider art cited therein during examination.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073.**
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 42624	Respectfully submitted, By:  Brian Siritzky, Ph.D. Registration No.: 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203
Main: (703) 894-6400 • FAX: (703) 894-6430

PTO/SB/08A (07-06)
Approved for use through 07/31/2006. OMB 0651-0031
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
JUL 12 2007
PATENT & TRADEMARK OFFICE

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
FORM PTO-1449 (modified)

Sheet 1 of 9

Application No.	11/065,412
Filing Date	Feb. 23, 2005
First Named Inventor	Farber, David A.
Group Art Unit	2157
Examiner Name	unassigned
Attorney Docket No.	2711-0073
Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	1-1	US-2004/0139097	15-Jul-04	Farber et al.
	1-2	US-2005/0038851	17-Feb-05	Kriegsman
	1-3	US-2005/0114296	26-May-05	Farber et al.
	1-4	US-4591983	27-May-86	Bennett et al.
	1-5	US-4594704	Jun-86	Ollivier
	1-6	US-4803641	7-Feb-89	Hardy et al.
	1-7	US-4847784	11-Jul-89	Clancey
	1-8	US-4943932	24-Jul-90	Lark et al.
	1-9	US-5029232	02-Jul-91	Nall
	1-10	US-5130792	14-Jul-92	Tindell et al.
	1-11	US-5132992	21-Jul-92	Yurt et al.
	1-12	US-5136716	08/1992	Churm et al.
	1-13	US-5191573	02-Mar-93	Hair
	1-14	US-5253275	12-Oct-93	Yurt et al.
	1-15	US-5675734	7-Oct-97	Hair
	1-16	US-5764906	06/1998	Edelstein et al.
	1-17	US-5802106	01-Sep-98	Packer
	1-18	US-5815664	09/1998	Asano
	1-19	US-5832514	11/1998	Norin et al.
	1-20	US-5845303	01-Dec-98	Templeman
	1-21	US-5893116	04/1999	Simmonds et al.
	1-22	US-5894554	05/1999	Lowery et al.
	1-23	US-5896533	04/1999	Ramos et al.
	1-24	US-5920701	07/1999	Miller et al.
	1-25	US-5935207	08/1999	Logue et al.
	1-26	US-5956489	09/1999	San Andres et al.

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 2 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	2-1	US-5966440	12-Oct-99	Hair
	2-2	US-5978791	02-Nov-99	Farber et al.
	2-3	US-5987606	11/1999	Cirasole et al.
	2-4	US-6002720	14-Dec-99	Yurt et al.
	2-5	US-6003030	14-Dec-99	Kenner et al.
	2-6	US-6014686	01/2000	Elnozahy et al.
	2-7	US-6018516	25-Jan-00	Packer
	2-8	US-6026440	02/2000	Sharder et al.
	2-9	US-6029176	22-Feb-00	Cannon
	2-10	US-6038216	14-Mar-00	Packer
	2-11	US-6038610	03/2000	Belfiore et al.
	2-12	US-6044405	28-Mar-00	Driscoll, III, et al.
	2-13	US-6046980	04-Apr-00	Packer
	2-14	US-6049831	11-Apr-00	Gardell et al.
	2-15	US-6065051	05/2000	Steele et al.
	2-16	US-6092204	07/2000	Baker
	2-17	US-6112239	29-Aug-00	Kenner et al.
	2-18	US-6115357	05-Sep-00	Packer et al.
	2-19	US-6119143	09/2000	Dias et al.
	2-20	US-6125388	26-Sep-00	Reisman
	2-21	US-6134583	10/2000	Herriot
	2-22	US-6144702	07-Nov-00	Yurt et al.
	2-23	US-6151624	11/2000	Teare et al.
	2-24	US-6154753	28-Nov-00	McFarland
	2-25	US-6167427	12/2000	Rabinovich et al.
	2-26	US-6173311	01/2001	Hassett et al.

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 3 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS				
Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	3-1	US-6181867	30-Jan-01	Kenner et al.
	3-2	US-6230196	05/2001	Guenthner et al.
	3-3	US-6243760	06/2001	Armbruster et al.
	3-4	US-6266699	07/2001	Sevcik
	3-5	US-6269394	22-Apr-03	Kenner et al.
	3-6	US-6282569	08/2001	Wallis et al.
	3-7	US-6298041	02-Oct-01	Packer
	3-8	US-6314565	06-Nov-01	Kenner et al.
	3-9	US-6347085	02/2002	Kelly
	3-10	US-6360256	03/2002	Lim
	3-11	US-6370571	04/2002	Medin, Jr.
	3-12	US-6370580	09-Apr-02	Kriegsman
	3-13	US-6412000	25-Jun-02	Riddle et al.
	3-14	US-6415280	02-Jul-02	Farber et al.
	3-15	US-6421726	16-Jul-02	Kenner et al.
	3-16	US-6456630	24-Sep-02	Packer et al.
	3-17	US-6480893	12-Nov-02	Kriegsman
	3-18	US-6484204	11/2002	Rabinovich
	3-19	US-6496856	17-Dec-02	Kenner et al.
	3-20	US-6502125	31-Dec-02	Kenner et al.
	3-21	US-6502215	12/2002	Raad et al.
	3-22	US-6557054	29-Apr-03	Reisman
	3-23	US-6581090	17-Jun-03	Lindbo et al.
	3-24	US-6591299	08-Jul-03	Riddle et al.
	3-25	US-6611862	26-Aug-03	Reisman
	3-26	US-6658464	02-Dec-03	Reisman

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 4 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

U.S. PATENT DOCUMENTS				
Examiner Initials*	Cite No.	Document No.	Publication/ Issue Date	Name of Patentee or Applicant of Cited Document
	4-1	US-6665706	16-Dec-03	Kenner et al.
	4-2	US-6741563	25-May-04	Packer
	4-3	US-6799221	28-Sep-04	Kenner et al.
	4-4	US-6915329	05-Jul-05	Kriegsman
	4-5	US-6928442	09-Aug-05	Farber et al.
	4-6	US-7047300	05/2006	Oehrke et al.
	4-7			
	4-8			
	4-9			
	4-10			
	4-11			
	4-12			
	4-13			
	4-14			
	4-15			
	4-16			
	4-17			
	4-18			
	4-19			
	4-20			
	4-21			
	4-22			
	4-23			
	4-24			
	4-25			
	4-26			

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 5 of 9		
	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.	Document No.	Publication Date	Name of Patentee or Applicant of Cited Document	T[1]
	5-1	EP-0801487 A2	10/15/1997	Law, et al.	
	5-2	EP-0817444	01/1998	Sun Microsystems	
	5-3	JP 10-027148	27/01/1998	Shuji, Ono et al.	
	5-4	JP-05162529	06/1993	Hiroshi, Saito	
	5-5	JP-08328583	09/1996	NT&T	
	5-6	JP7066829	03/1995	Hiroshi, Saito	
	5-7				
	5-8				
	5-9				
	5-10				
	5-11				
	5-12				
	5-13				
	5-14				
	5-15				
	5-16				
	5-17				
	5-18				
	5-19				
	5-20				
	5-21				
	5-22				
	5-23				
	5-24				
	5-25				
	5-26				

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 6 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES			
Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
	6-1	"Exporting Web Server Final Report," http://www.cs.technion.ac.il/Labs/Lccn/projects/spring97/project4/final_report.html, Spring 1997 (downloaded 07/07/07).	
	6-2	Andresen, D., et al., Multiprocessor scheduling with client resources to improve the response time of WWW applications, Proc. 11th Int'l Conf. on Supercomputing (Austria, July, 1997). ICS '97. ACM Press, NY, NY, 92-99.	
	6-3	Basturk, E., et al., "Using network layer anycast for load distribution in the Internet," Tech. Rep., IBM T.J. Watson Research Center, 1997 (21 pgs.).	
	6-4	Berners-Lee et al., RFC 1738--Uniform Resource Locators, Dec. 1994.	
	6-5	Bestavros, et al., "Server-Initiated Document Dissemination for the WWW," IEEE Data Engineering Bulletin 19(3):3-11, Sept. 1996.	
	6-6	Bhattacharjee et al., "Application-layer anycasting," in Proc. IEEE INFOCOM '97, 1997.	
	6-7	Braun, H., et al., "Web traffic characterization: an assessment of the impact of caching documents from NCSA's web server," Comput. Netw. ISDN Syst. 28, 1-2 (Dec. 1995), 37-51.	
	6-8	Brisco, T. P. RFC 1794: DNS support for load balancing, Apr. 1995.	
	6-9	Carter et al., "Dynamic server selection using bandwidth probing in wide-area networks," Tech. Rep. BU-CS-96-007, Comp. Sci. Dept., Boston University, Mar. 1996.	
	6-10	Carter et al., Server selection using dynamic path characterization in Wide-Area Networks, IEEE INFOCOM '97, 1997 (pp. 1014-1021).	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 7 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES

Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
	7-1	Cisco Systems, Inc., Cisco DistributedDirector 2500 Series Installation and Configuration Guide, pp. xix - xxii; 1-1 to 1-12; 6-1 to 6-18; 7-1 to 7-18; 8-1 to 8-24, pub. date unknown, downloaded 04/07: http://www.cisco.com/univercd/cc/td/doc/product/iaabu/distrdir/dd2501/.	
	7-2	Cisco Systems, Inc., Cisco DistributedDirector 4700-M Installation and Configuration Guide, pgs. xix-xxii; 1-1 to 1-14; 7-1 to 7-18, 8-1 to 8-20; pub. date unknown, [downloaded 04/07 from http://www.cisco.com/univercd/cc/td/doc/product/iaabu/distrdir/dd4700m/].	
	7-3	Cohen, J., et al., "Cache Array Routing Protocol v1.1", 29 Sept. 1997; http://tools.ietf.org/id/draft-vinod-carp-v1-01.txt (Last-Modified: Wed, 01 Oct 1997).	
	7-4	Colajanni, M. and Yu, P. S. 1997. Adaptive TTL schemes for load balancing of distributed Web servers. SIGMETRICS Perform. Eval. Rev. 25, 2 (Sep. 1997), 36-42.	
	7-5	Crovella et al., Dynamic server selection in the Internet, 3rd IEEE Workshop on the Arch. and Implementation of High Performance Computer Sys. '95, pp. 158-162, Aug. 1995.	
	7-6	Danzig, P. B., et al., "An analysis of wide-area name server traffic: a study of the Internet Domain Name System," Conf. Proc. Communications Architectures & Protocols (Aug. 1992). D. Oran, Ed. SIGCOMM '92. ACM Press, New York, NY, 281-292.	
	7-7	Doi, K. "Super Proxy Script--How to make distributed proxy servers by URL hashing," Sharp Corp., http://naragw.sharp.co.jp/sps/, dates unknown (1996 -2000), download 07/07/07.	
	7-8	Fox, A., "A Framework for Separating Server Scalability and Availability from Internet Application Functionality," PhD thesis, University of California, Berkeley, 1998.	
	7-9	Fox, A., et al,, "Cluster-based scalable network services", Proc. 16th ACM Symp. on Operating Systems Principles (Saint Malo, France, Oct. 05 - 08, 1997). W. M. Waite, Ed. SOSP '97. ACM Press, New York, NY, 78-91.	
	7-10	Goldszmidt, et al., "Load Distribution for Scalable Web Servers: Summer Olympics 1996--A Case Study," In Proc. 8th IFIP/IEEE Int'l Workshop on Distributed Systems: Operations and Management, Sydney, Australia, Oct. 1997.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 8 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES			
Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
	8-1	Gulbrandsen, A., et al., "A DNS RR for specifying the location of services (DNS SRV)", Network Working Group, RFC 2052, Oct. 1996.	
	8-2	Guyton et al., "Locating nearby copies of replicated Internet servers," Proc. ACM SIGCOMM '95, pp. 288-298, 1995.	
	8-3	Jeffrey et al., Proxy-Sharing Proxy Servers, IEEE, pp. 116-119, 1996.	
	8-4	Kwan et al., NCSA's World Wide Web Server: Design and Performance, IEEE, pp. 68-74, Nov. 1995.	
	8-5	Luotonen et al., World-Wide Web Proxies, CERN, Apr. 1994 (modified May 24, 1994).	
	8-6	Mockapetris et al., "Development of the Domain Name System," Proc. SIGCOMM '88 Computer Communications Review, vol. 18, No. 4, Aug. 1988.	
	8-7	Mockapetris, P., RFC 1034: Domain Names - Concepts and Facilities, Nov. 1987.	
	8-8	Mockapetris, P., RFC 1035: Domain Names - Implementation and Specification, Nov. 1987.	
	8-9	Oguchi et al., A Study of Caching Proxy Mechanisms Realized on Wide Area Distributed Networks, High Performance Distributed Computing, 5th Int'l Symposium, pp. 443-449, 1996.	
	8-10	Peleg D., et al., "The Availability of Quorum Systems," Tech. Rep. CS93-17, The Weizmann Institute of Science, Rehovot, Israel, 1993.	

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT FORM PTO-1449 (modified) Sheet 9 of 9	Application No.	11/065,412
	Filing Date	Feb. 23, 2005
	First Named Inventor	Farber, David A.
	Group Art Unit	2157
	Examiner Name	unassigned
	Attorney Docket No.	2711-0073
	Confirmation No.	3296

NON-PATENT REFERENCES			
Examiner Initials*	Cite No.	Non-patent Reference bibliographic information, where available	T[1]
	9-1	Petri S., et al., "Load Balancing and Fault Tolerance in Workstation Clusters. Migrating Groups of Communicating Processes.", Operating Systems Review, Vol. 29, No. 4, Oct.1995, pp. 25-36.	
	9-2	Postel, RFC 1591--Domain Name System Structure and Delegation, Mar. 1994.	
	9-3	Ross, K.W., "Hash-Routing for Collections of Shared Web Caches", IEEE Network Magazine, 11, 7:37-44, Nov-Dec 1997.	
	9-4	Schemers, R., "lbnamed - A load balancing name server written in Perl," 1995 LISA IX – Sept. 17-22, 1995 – Monterey, CA.	
	9-5	Schuba, Christoph; "Addressing Weaknesses in the Domain Name System Protocol," COAST Laboratory, Dept. of Computer Sciences, Purdue University; West Layfayette, IN; Aug. 1993, p. 1-87.	
	9-6	Smith, "What can Archives offer the World Wide Web?", Technical Report 11, University of Kent, Computing Laboratory, University of Kent, Canterbury, UK, 1994.	
	9-7	Thaler, D. G. and Ravishankar, C. V. 1998. Using name-based mappings to increase hit rates. IEEE/ACM Trans. Netw. 6, 1 (Feb. 1998), 1-14.	
	9-8	Vixie, Paul; "Name Server Operations Guide for BIND," Internet Software Consortium; La Honda, CA; p. SMM:10-2--SMM:10-30 (undated, 1996).	
	9-9	Wessels, Intelligent Caching for World-Wide Web Objects, Masters Thesis, University of Colorado, 1995.	
	9-10		

Examiner Signature		Date Considered	

*Examiner: Initial if reference was considered, whether or not citation is in conformance with MPEP 609. Draw a line through citation if not in conformance and not considered. Include a copy of this form with next communication to applicant.

[1]Translation provided.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 817 444 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.01.1998 Bulletin 1998/02

(51) Int Cl.[6]: **H04L 29/06**

(21) Application number: **97304697.2**

(22) Date of filing: **30.06.1997**

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE

(30) Priority: **01.07.1996 US 674561**

(71) Applicant: **SUN MICROSYSTEMS, INC.**
Mountain View, California 94043-1100 (US)

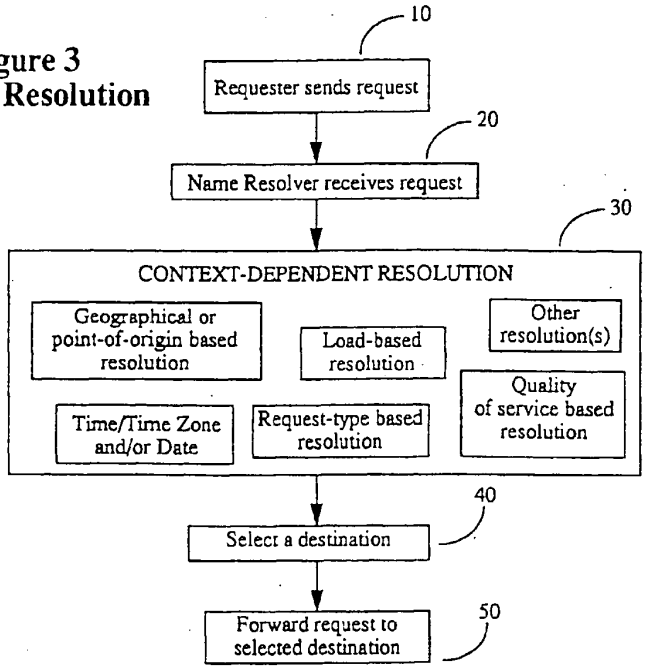
(72) Inventor: **Ebrahim, Zahir**
Mountain View, California 94043 (US)

(74) Representative: **Turner, James Arthur et al**
D. Young & Co.,
21 New Fetter Lane
London EC4A 1DA (GB)

(54) System for context-dependent name resolution

(57) A context-dependent, multiply binding name resolution system. A name resolver is provided, connected to either a requester's system or a receiver's system, or both. Requests to a given service or domain name are resolved to the appropriate IP address. The intended recipient of the request is resolved based upon a combination of one or more predetermined criteria, including: information about the sender (e.g. geographical location, specific requester identity, etc.); information about the intended recipient (e.g. load balance at the receiver, type of service, etc.); information contained within the request itself (e.g. type of service requested); or other information (time of day, date, random selection of recipient, e.g.). The system is implemented in hardware and/or software, and the resolution criteria can be made interdependent or independent.

Figure 3
Service Resolution



EP 0 817 444 A2

Description

This invention relates to systems for context-dependent name resolution.

In a network (either Intranet or Internet) system, when a host wants to communicate with another host or locate a service or another host on the network, the address is typically resolved in an absolute manner; that is, the naming service of the network resolves a given name to a single IP (Internet Protocol) address. For instance, a domain name service (DNS) call resolves a name to a single IP address or host name on the Internet.

When there are many connections to a given host or service, the resulting congestion degrades the overall performance. The service provider may upgrade the service, e.g. to increase its capacity. In this case, in order to keep the modification of the service transparent to requesting hosts, a demultiplexer may be used, as illustrated by way of example in Figures 1-2. For instance, in Figure 1 a call to Service 3 will be uniquely resolved to that service's IP address and transmitted there accordingly, because there is only a single binding in the DNS name resolver to the host of Service 3. In general, in a conventional system such as that illustrated in Figures 1, 2 and 2A, there will be a single binding between a name and an IP address.

Yet another alternative is for a special type of demultiplexer, which may monitor the traffic, and from the different IP addresses can collect statistics such as response time and/or the load on the corresponding IP hosts, and use this information for load balancing by binding a requested name to an IP address for a host which is less heavily loaded.

If Service 3 is upgraded to include Services 3.1 through 3.Y, then the demultiplexer is added (see Figure 2); the requesting or DNS resolving host resolves the request as usual and sends it out over the (Inter- or Intra-)network, but the demultiplexer is interposed between the network and the service(s), to route the incoming requests to one of the services. Note that in this case, there is still a single binding between the DNS and the demultiplexer.

A problem with this approach is that the demultiplexer acts as a limiting factor on throughput, and additionally represents a single source of failure for all of the services that it serves. A mechanism is needed whereby such a critical point and bottleneck can be avoided, while still providing access to the services for remote users.

An alternative approach to decreasing congestion is through DNS spoofing, where the name resolver is fooled into giving out a different name binding, by keeping the name resolver updated frequently with a new name resolution binding based on some criteria such as load distribution on target servers. An example of DNS spoofing is shown in Figure 2A, in which a requester 2 issues a request using the name "sun.com", and a DNS lookup 4 looks it up. The host (1, 2 or 3) to which the names is bound will depend upon criteria employed at the destination domain, as determined by DNS spoofing service 6. Such criteria may include load on the hosts, their availability, etc. The DNS spoofing service 6 polls the hosts 1-3 periodically or at demand, and binds one of them at a time to the DNS lookup name resolver for the name "sun.com".

In a spoofing service, the criteria used to achieve the binding are not related to the context of the requester; they rely only upon information at the destination. Thus, much information that could be important in making a binding is not utilized. In none of the known systems is caller context used.

It would accordingly be advantageous to have a system which takes into account the context of the requester in the process of name resolution.

Various respective aspects and features of the invention are defined in the appended claims.

A context-dependent name resolution system is presented, which implements predetermined criteria for static or dynamic binding of a "name" to any one of several "objects" of the same type. Which "object" is selected for the binding is determined by a "name resolver" using context information that is explicitly specified by the requestor in a name resolution lookup call to the "name resolver". The "name resolver", which is implemented as a server program, stores a lookup table comprising bindings of a "name" to several different instances of an object, along with "policy" information which defines the criteria for selecting an instance of the object. For instance, the name may be an internet URL of the form "www.sun.com", and the resolved to object may be an IP address of the type 192.146.85.82. Thus the "name resolver" will store the tuple: <www.sun.com> <IP_address_1> <IP_address_2>..., and will store policies on which IP address to bind to www.sun.com when a lookup request specifying "www.sun.com" is received by it.

By enabling multiple binding support in the "name resolver", several advantages are realized. First, the system enables new resources to be added transparently to the caller. When a new server is added to distribute the load for one or more hosts corresponding to the name "sun.com", for instance, the IP address for the new server is simply added to the tuple in the name resolver, along with any desired policy criteria, e.g. to use only between 9 a.m. and 5 p.m.

Secondly, the fundamental limit to scaling which is inherent in today's single binding systems is removed, by enabling multiple destination servers providing the same service to be referenced by the same "name" handle by the caller (for example "sun.com"), and the name to be resolved by the name resolver to different physical computers servers, depending upon caller context. Many instances of the name resolver may be running on different physical computers, and many instances of the services may be running, and the requester will ac-

cess the name resolver closest to him; and the resolved object that is closest to the requester or that can best serve the request (based upon some selected or predetermined criteria) will be bound to the request, all transparently to the requester. Another advantage of the system of the invention is that it allows resources to be easily replicated transparently to the user, eliminating any single point of failure in the network or at the service end point. If one valid destination fails, is not responding or is overloaded, the name resolver is already capable of binding a "name" to multiple destinations, and the unavailable destination can be disabled in the name resolver's internal tables, using administrative means.

The context that may be used as a basis for the binding is variable; it may include information about the requester (e.g. requester IP address, domain name or inferred geographical region), about the destination (with similar alternatives), about the request itself (e.g. type or quality of service requested), or about independent information (e.g. time of day or time zone of the point of origin of the request).

The system may be implemented (eg as part of a computer apparatus) using hardware, software, or combinations thereof.

The invention will now be described by way of example with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

Figures 1-2A are diagrams of present network systems used for destination addresses for requests.

Figure 3 is a flow chart illustrating a method according to an embodiment of the present invention.

Figures 4 and 5 are diagrams of network systems incorporating features of embodiments of the present invention.

Figure 6 in an illustration of the operation of an embodiment of the invention using geography-based resolution criteria.

Figure 7 is a diagram depicting zones in which domain name service resolvers and service location resolvers of the invention may be located.

A method according to the invention is illustrated in Figure 3, and Figures 4-5 shows a suitable systems implementable on the Internet or WorldWide Web, the Internet or an Intranet,, incorporating features of the name resolution system.

Name resolution can include any kind of name resolution lookup for binding one object to another. This includes binding the name of a service to a host computer (or its IP address) that provides that service, and binding the type of service to the name of a service providing that type of service. It also includes resolving a name in one domain to another name in the same domain, and to another name in a different domain. For the purposes of this invention, the terms "service" and "host" can be considered synonymous, as can "service location" and "host location".

The system includes three fundamental features that can be implemented singly or in any combination with one another: multiple binding of destination addresses to names; caller context name resolution in connection with such multiple binding; and name resolution policy considerations, i.e. independent criteria upon which such resolution is based.

Name Resolution : An Example

An example of name resolution occurs, for instance, when a user wishes to view video, e.g.a movie, on his or her computer. An application running on the computer is used to make a procedure call to the name resolver to determine which server can serve a movie, e.g.

service_handle = name_service_lookup("movie")

where the string "movie" is passed to a server, which then returns the name of the service providing that movie (or it may return a linked list of services providing movie services, from which the user may be prompted to select one).

The server may, by way of example, return a single service name (e.g. "quicktime") in the variable called "service_handle". Next, the application must determine which server host can provide this service to this client; to do this, it makes another name resolution call:

host_handle = name_hostname_lookup (service_handle)

and passes it the variable "service_handle" returned by the previous call. This call returns it the name of a server which is listed as offering this service.

The network location of this server must now be determined, so that the user's computer or workstation can connect to it to get the movie service, and so it makes the call:

host_IP_address = nome_hostaddress_lookup (host_handle)

and finally obtains an IP address of an appropriate host that can provide the movie service to the user.

Further calls to this quicktime server may determine a list of movies that are available. Alternatively, a design may be had in which the caller simply specifies the name of the movie he wants to see, and the name lookup returns to the caller the IP address of a host that is currently able to serve the requested movie to the caller.

All of or some of the above steps may be combined in more integrated calls to the name resolver.

In the present invention, in addition to the above, an additional parameter is passed to the name_hostaddress_lookup() function above (or to all the functions above); this parameter may be called the caller_context. This caller_context is used by the name resolver to help it decide which IP address to use. Thus, the calls would look like:

host_IP_address = name_hostaddress_lookup (host_handle, caller_context). Caller_context is a cookie (or structure) that may include information such as the caller's IP address, its point of origin, the quality of the requested service, or any other context that might be

relevant. These contexts must be well specified and their format(s) standardized, so that many different name resolvers can interoperate properly. The particular caller_context formats and contents selected are not themselves crucial,

The present embodiment uses requester context-dependent intelligence to determine which of several destination hosts should be utilized, or in other words, which of several IP addresses should be used to resolve a given "name" request. The context-dependent intelligence, implemented as hardware or software or a combination of both, allows the "name" resolution to be based upon some context information from the requestor, and/or upon the "name" being resolved in a context-free manner, but optionally dependent upon criteria such as load balancing on the destination.

As shown in Figure 4, a requesters 100-120, of which there may be an arbitrary number, each comprise computers or workstations on a local network, served by a server 150. The requesters will typically be conventional workstations, personal computers, or any networkable computing device, with local processors (130, etc.), memory (140, etc.), mass storage (not shown) and network connections. The server 150 also includes a conventional processor 160, memory 170, mass storage and necessary connections and control hardware and software.

A name resolver 180 forms part of the server 150, or may be a separate unit. It may be implemented entirely in software, i.e. as program modules stored in mass storage and/or the memory 170, or it may be a stand-alone device, with its own logic or processor and memory, as desired.

The server 150 is connected to a broader network 170, such as the Internet, an Intranet or WAN (wide-area network), or the WorldWide Web, via a network connection (e.g. a T-line) 260. On this network 170 is optionally another name resolver 200, which may be implemented in the same fashion or differently from the name resolver 180. Name resolvers 180 and 200 may both be used, or either may be used by itself. (See discussion of Figure 7, below.)

Destination hosts 210-230 may also be conventional computers or workstations with processors (such as 240), memory (such as 250), mass storage (not shown), network connections, etc., as needed to implement conventional network functions in addition to the features of the present invention.

Referring to the flow chart of Figure 3, when a requester (such as 100-120) sends a name resolution request to a name resolver 180, instead of the single binding of conventional systems, the name resolver 180 includes multiple binding tables and/or functions (implemented by appropriate logic or program modules) to resolve the request to an appropriate IP address for a destination host (such as 210-230 in Figure 4).

An example of such a multiple binding would be the binding of multiple, geographically disparate destinations to a single domain name. If, for instance, a user in Germany wishes to access www.sun.com (the WWW server for Sun Microsystems, Inc.), he or she can simply use the Uniform Resource Locator (URL) "http://www.sun.com"). Thus, the user sends the request from requester 310 (see Figure 6), and the name resolver 320 determines that the requester is in Germany. (See step 20 in Figure 3). The selected destination may either be in the United States or elsewhere in Europe, since Sun Microsystems in this example has two "sun.com" locations. Since there is a valid European local destination 330, the name resolver 320 resolves the destination address to sun.com (Europe), as indicated at box 30 of Figure 3. This resolved destination is then selected for the request packet(s) (box 40 of Figure 3), and the request is forwarded to sun.com (Europe) (box 50 of Figure 3).

Following are lists of context-dependent resolution criteria applicable to the invention and usable at box 30 of Figure 3, where resolution may be based upon requester information, destination information, request contents, or other information:

A. Resolution Based Upon Requester Information

Examples of manners in which the resolution dependent upon requester information may be carried out include:

1. resolution based upon the domain name of the sender;
2. resolution based upon the inferred (looked-up) actual geographic region of the sender;
3. resolution based upon other geography-related information of the sender:

    4. either directly ascertainable from the request (city/country info, e.g.); and/or
    5. indirectly ascertainable (area code, address, etc.);

6. resolution based upon quality of service desired by the requester; and
7. resolution based upon requester's time of day or time zone.

B. Resolution Based Upon Destination Information

Examples of manners in which the resolution dependent upon a specified or intended destination or receiver information may be carried out include:

1. resolution based upon the load at the receiver;
2. resolution based upon the inferred (looked-up) actual geographic region of the receiver;
3. resolution based upon other geography-related information of the receiver:

4. either directly ascertainable from the request (city/country info, e.g.); and/or
5. indirectly ascertainable (area code, address, etc.).

Criterion B.1 can be resolved either based upon information as perceived at the sender end or based upon information as it exists at the receiver end, and resolved at the sender end. For example, the sender might frequently send requests to a particular server, service, geographical region or organization. In this case, it may be advantageous to regularly poll one or more often-used destinations to determine their level(s) of activity. Alternatively, such polling could automatically be carried out for any destinations for which the sender determines that its number of requests exceeds a certain threshold, or for a top predetermined percentage of requests sent by the requester (e.g. all destinations for which the number of requests exceeds N1%, or the top N2 destinations to which the sender makes requests, where N1 and N2 are appropriate predetermined numbers). Polling can be implemented as program modules integral to the name resolver(s), either entirely in one or in more than one location (e.g. partially in a local name resolver and partially in the destination server).

Alternatively or in addition, criterion B.1 may be based upon information at the receive end at the time of sending the request, and resolved at the receiver.

C. Resolution Based Upon Request Contents

1. type of service requested;
2. specific information (or type of information) requested;
3. any other implicit information inferred from the request; and/or
4. any other explicit information obtained from the request.

D. Resolution Dependent Upon Other Factors

1. randomly generated selection of destination based upon qualified list; and/or
2. other independently developed information.

These resolution criteria can be used singly or in any desired combination with one another, as specified by the requester or the administrator of the destination (e.g. server owner or service provider). For instance, in the example of Figure 6 the resolution of the destination address to sun.com (Europe) could be modified if the load of sun.com (Europe) is larger by some predetermined threshold amount than the load at sun.com (USA) -- in which case the sun.com (USA) resolution could override the default sun.com (Europe) resolution for requesters in Europe, either at all times or only during peak European hours, which correspond to off-peak hours in the United States.

An extension of this system is in the resolution of services provided at the receiver network of a request. For instance, a given organization may have many database servers, and request for access to given information may in conventional systems be routed to the same server, because of the single binding nature of destination resolution. With multiple address resolution binding, multiple database servers (or other destination hosts) such as 210, 270, 280, etc. (see Figure 5) may be made available under a single name, in which case a local service resolution server, such as server 300 in Figure 5, is provided. The server 300 includes tables, program modules or the like as desired to resolve a request to any one of several to many IP addresses. The selection of a given server to receive a request can depend upon, for example, load at each of the servers.

In general, the resolution systems or subsystems will be in one of three "zones" (see Figure 7): the sender's zone 1 (up to, e.g., the sender's firewall or Internet server); the name resolver system's zone 2 (which may be at a server on the Internet, or any place outside zone 1 which is not in the receiver's system; and the receiver's zone 3. Service location name resolvers may be in any of zones 1-3; typically, destination lookup name resolvers, which may comprise multiple-binding domain name services (DNS's), will be in either zone 1 or zone 2. It will be understood that zones, for the purpose of this application, are administrative boundaries, and need not correlate to actual network boundaries.

The present embodiment can interact with existing systems in a number of ways. For instance, in secure systems where the source address is stripped from the request by a firewall (and only, e.g., the domain name remains), then the entire source address would not be used to resolve the destination; the name resolver is provided with the program instructions or circuitry to implement this. Note also that the DNS name resolution may be a distinct operation from service location, and thus the destination can be selected due to a combination of considerations including source address, requested destination address, and request contents, including the type of request made or service requested.

The multiple binding feature of the invention allows a single name, Internet address (e.g. URL address such as http://www.sun.com) or the like to represent as many actual servers (or services) as desired. This has a distinct advantage of allowing modifications, upgrades or replications to be made to the destination server(s) without modification of the destination address or service name as used by the requester. Such modifications may include adding servers or services, or establishing a mirror site at a location geographically near a large source of requests, etc.

There may be security and consistency implications of multiple bindings, with concomitant solutions. If a destination address can be multiply bound, a name resolver may be configured to bind to an incorrect address, either deliberately or otherwise. While this may also be a prob-

lem with existing single binding, multiple binding may complicate detection of the problem. Thus, the system of the invention may be combined with programs and tables for verification that a given binding is valid. Very secure systems may wish to limit their multiple binding to a predetermined, fixed number of resolution possibilities, or even maintain the current system of single binding as far as Internet or Web transmission is concerned, and implement multiple binding only behind a firewall. Alternatively or in addition, encryption and digital signatures may be used to ensure validity of bindings, and of modifications or additions to the bindings.

Variations may be had by designating a given name resolution as a default resolution, and utilizing alternative resolutions only upon the request meeting predetermined criteria (such as load imbalance, geographical distance, specific sources of the request, etc.).

Additional bindings (i.e. additional addresses to which a request may resolve) may be added without the requester being aware of it, and indeed substitutions may be made at any time, again transparently to the requester.

F. Service Selection Based Upon Caller Context

The system of the invention may be applied more generally to the use of caller context for name resolution. Context about the requesting user (e.g. in a variable called "caller_context") can be passed to appropriate functions, such as:

service_handle = name_service_lookup("movie", caller_context)

host_handle = name_hostname_lookup (service_handle, caller_context)

Here, the function service_handle executes a service name lookup based upon the input "movie", as well as the information in the variable caller_context, outputting the value service_handle. Given this output as input, along with (again) caller_context, the function host_handle then returns an appropriate host name.

An example of a possible such situation could result when a user desires an encryption service, and several encryption policies are available. The service_handle function can be configured to return different encryption algorithm services based upon the specified destination (which is specified in caller_context), or there may be an effective override in the caller_context by which the user can select a higher level of security.

In general, as noted above, the invention may be implemented at the sender's system or at the destination system, or indeed at points in between (e.g. at proxy servers). Note that single points of failure and bottlenecks are removed using the present system, and a truly distributed, multiply binding name resolution system is achieved.

Particular and preferred aspects of the invention are set out in the accompanying independent and dependent claims. Features of the dependent claims may be combined with those of the independent claims as appropriate and in combinations other than those explicitly set out in the claims.

Claims

1. A system for name resolution of at least one request directed to a destination name of a receiver and issued from a requester system, including:

   a name resolver connected to the requester service and including

   predetermined information correlating a plurality of destination addresses with said destination name;

   a criteria selection subsystem configured to provide a selection of at least one said destination address based upon at least one predetermined criterion; and

   a destination address binder configured to bind said destination name with said selected destination address.

2. The system of claim 1, wherein said at least one predetermined criterion includes a criterion based upon information explicitly contained within said request.

3. The system of claim 2, wherein said information includes an address of said requester.

4. The system of claim 2, wherein said information is based upon at least a portion of contents of said request.

5. The system of claim 4, wherein said portion of said contents includes at least one of:

   a type of service requested by said request; and

   a quality of service requested by said request.

6. The system of claim 1, wherein said at least one predetermined criterion includes a criterion based upon information not explicitly contained within said request.

7. The system of claim 6, wherein said information pertains to a geographical region of said requester.

8. The system of claim 6, wherein said information pertains to a geographical region of said receiver.

9. The system of claim 6, wherein said information includes load information relating to said receiver.

10. The system of claim 6, wherein said information includes at least one of time and date information.

11. The system of claim 1, wherein said binding is to a port at said destination address.

12. The system of claim 1, wherein said binding is based at least in part upon a random selection of one said destination address.

13. A method for context-dependent binding of a destination name to a destination address in a name resolver having a predetermined set of binding criteria and predetermined correlations of destination names with destination addresses, wherein at least one said correlation is a multiple correlation correlating a plurality of said destination addresses with a single destination name, including the steps of:

(1) receiving said request at said name resolver; and
(2) based upon said multiple correlation, retrieving at least one said destination address from said plurality of destination addresses.

14. The method of claim 13, further including after step 2, the step of:

(3) binding said retrieved destination address to said destination name.

15. The method of claim 14, further including after step 3, the step of:

(4) transmitting said request with said destination address.

16. The method of claim 13, wherein said multiple correlation is based upon information explicitly contained within said request.

17. The system of claim 13, wherein said multiple correlation is based upon at least a portion of contents of said request.

18. The system of claim 13, wherein said multiple correlation is based upon at least one of:

an address of said requester; and
a type of service requested by said request; and
a quality of service requested by said request.

19. The system of claim 1, wherein said multiple correlation is based upon information not explicitly contained within said request.

20. The system of claim 19, wherein said multiple correlation is based upon at least one of:

a geographical region of said receiver;
load information relating to said receiver;
time information;
date information; and
a random factor generated by said name resolver.

21. A system for name resolution of at least one request directed to a destination name of a receiver and issued from a requester system, the request information relating to caller context, the system including:

a name resolver connected to the requester service and including

predetermined information correlating a plurality of service names with said caller context;
a service selection subsystem configured to provide a selection of at least one said service name based upon said caller context; and
a host name selection subsystem configured to return a host name based corresponding to said selected service name.

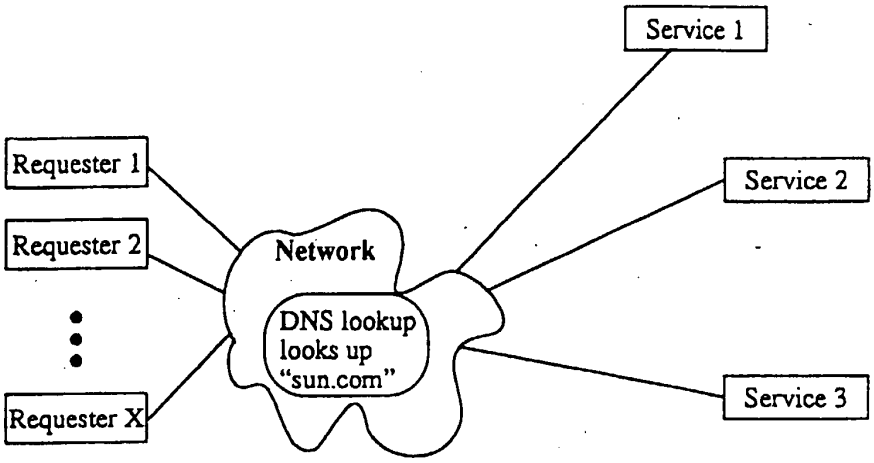
22. The system of claim 1, wherein said at least one predetermined criterion includes:

a first criterion based upon information explicitly contained within said request; and
a second criterion based upon information not explicitly contained within said request.

23. The system of claim 22, wherein said second criterion corresponds to a predetermined policy used by said name resolver.

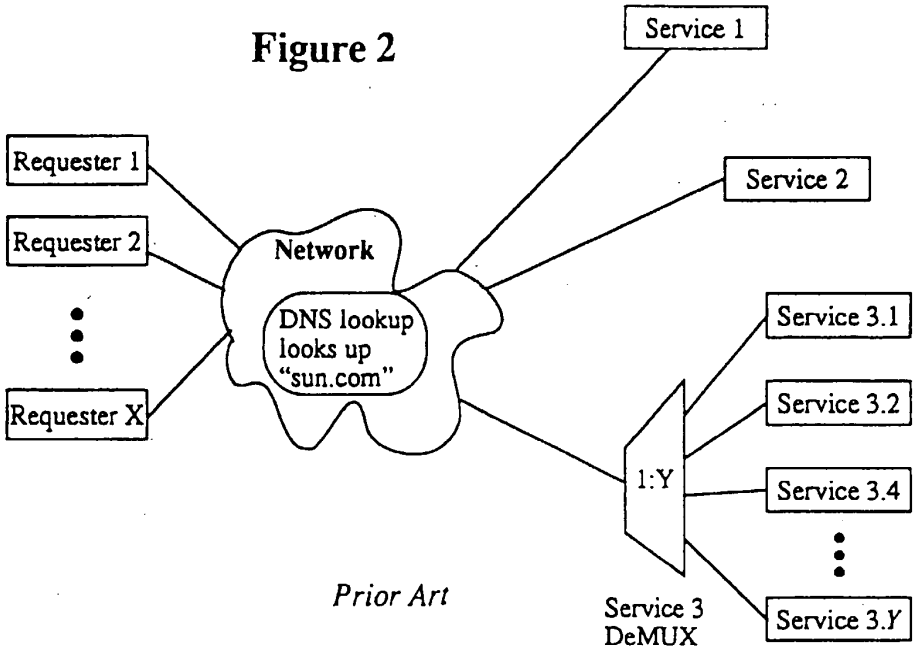
24. The system of claim 23, wherein said first criterion corresponds to information relating specifically to said requester system.

Figure 1



Prior Art

Figure 2



Prior Art

Figure 2A

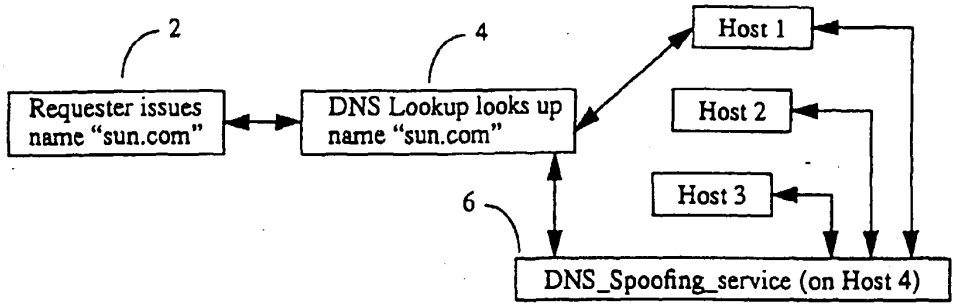


Figure 3
Service Resolution

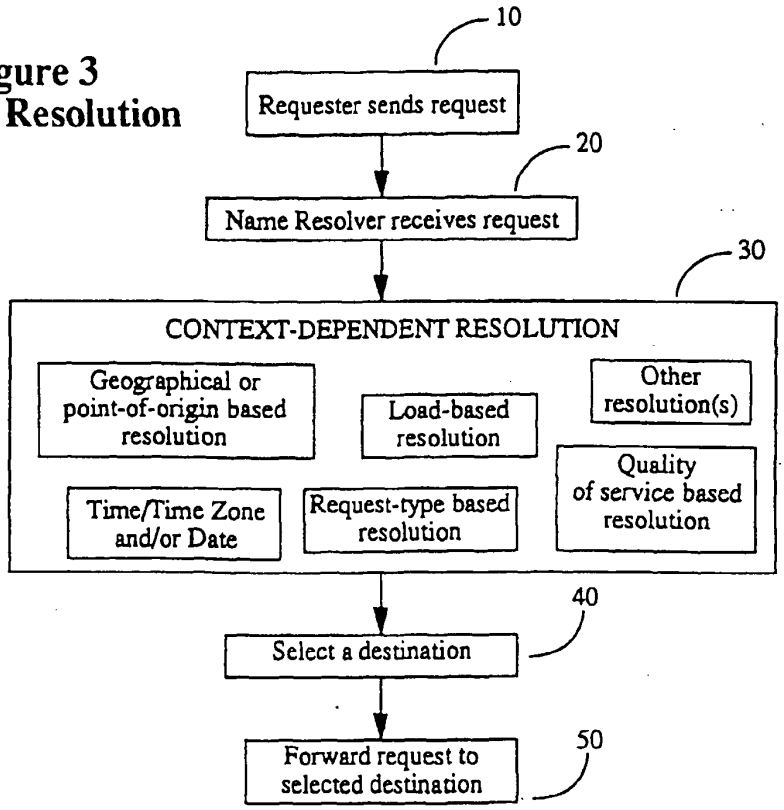


Figure 4

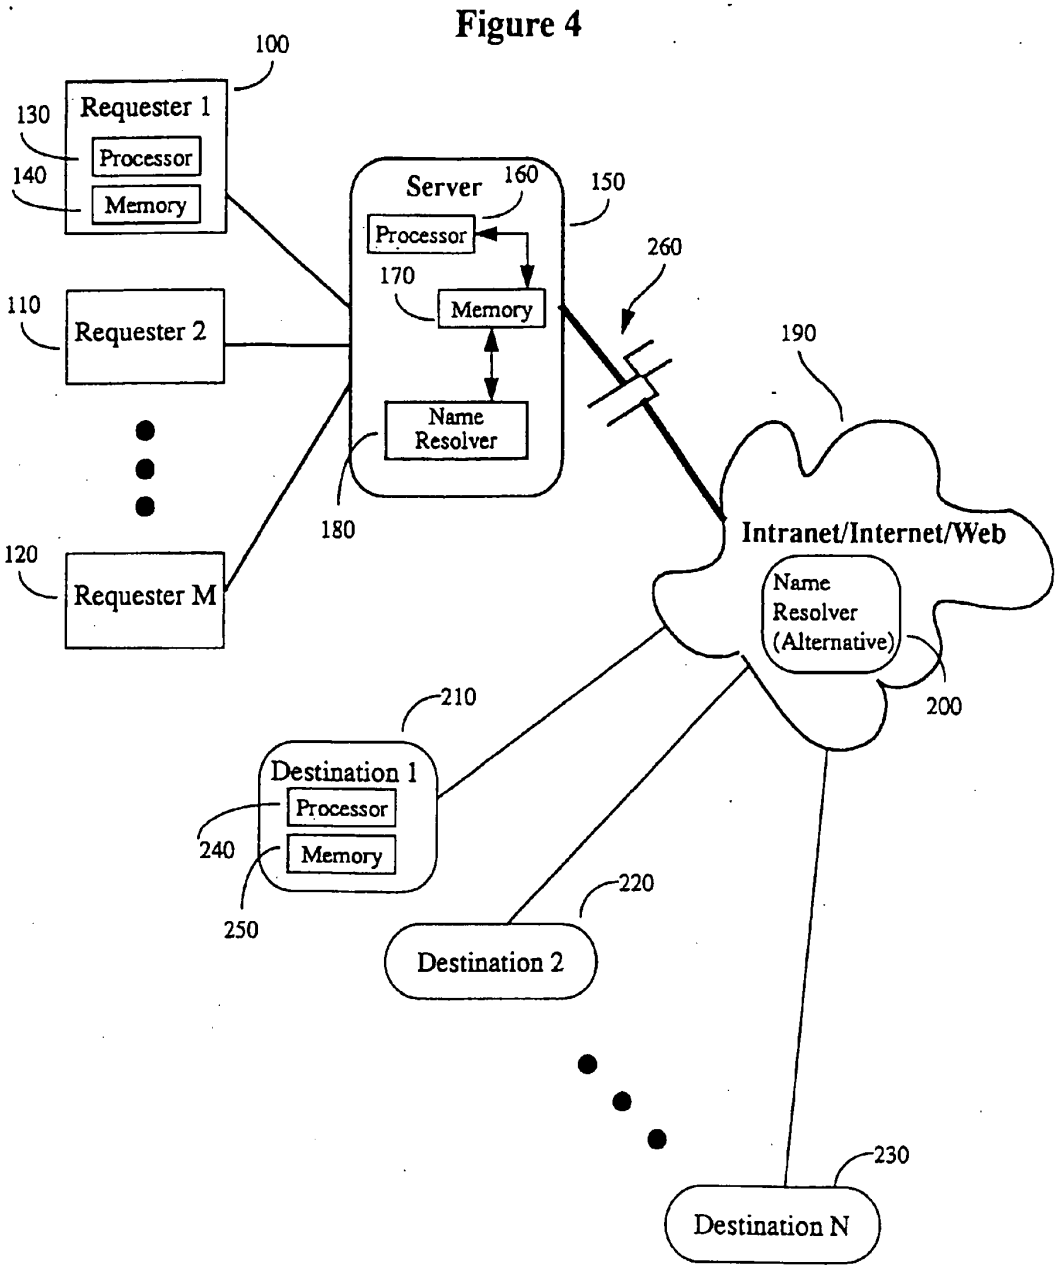

100
Requester 1
130
Processor
140
Memory
Server
160
150
Processor
260
170
Memory
110
Requester 2
190
Name
Resolver
180
Intranet/Internet/Web
120
Requester M
Name
Resolver
(Alternative)
200
210
Destination 1
Processor
240
Memory
250
220
Destination 2
230
Destination N

Figure 5

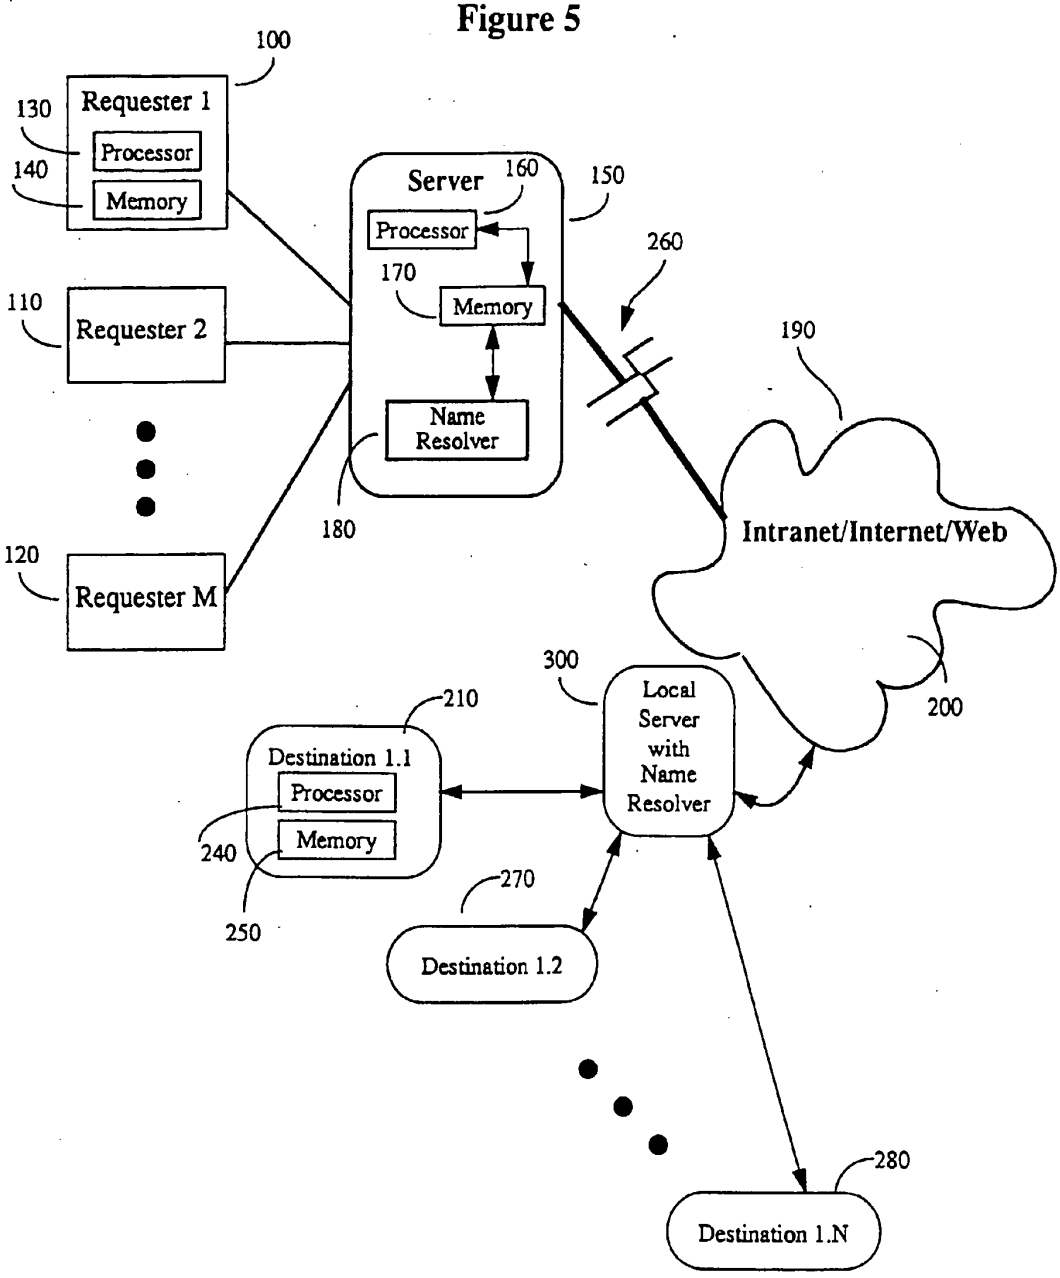

100
Requester 1
130
Processor
140
Memory
Server
160
150
Processor
260
170
Memory
110
Requester 2
190
Name
Resolver
180
Intranet/Internet/Web
120
Requester M
300
Local
Server
with
Name
Resolver
200
210
Destination 1.1
Processor
240
Memory
250
270
Destination 1.2
280
Destination 1.N

Figure 6

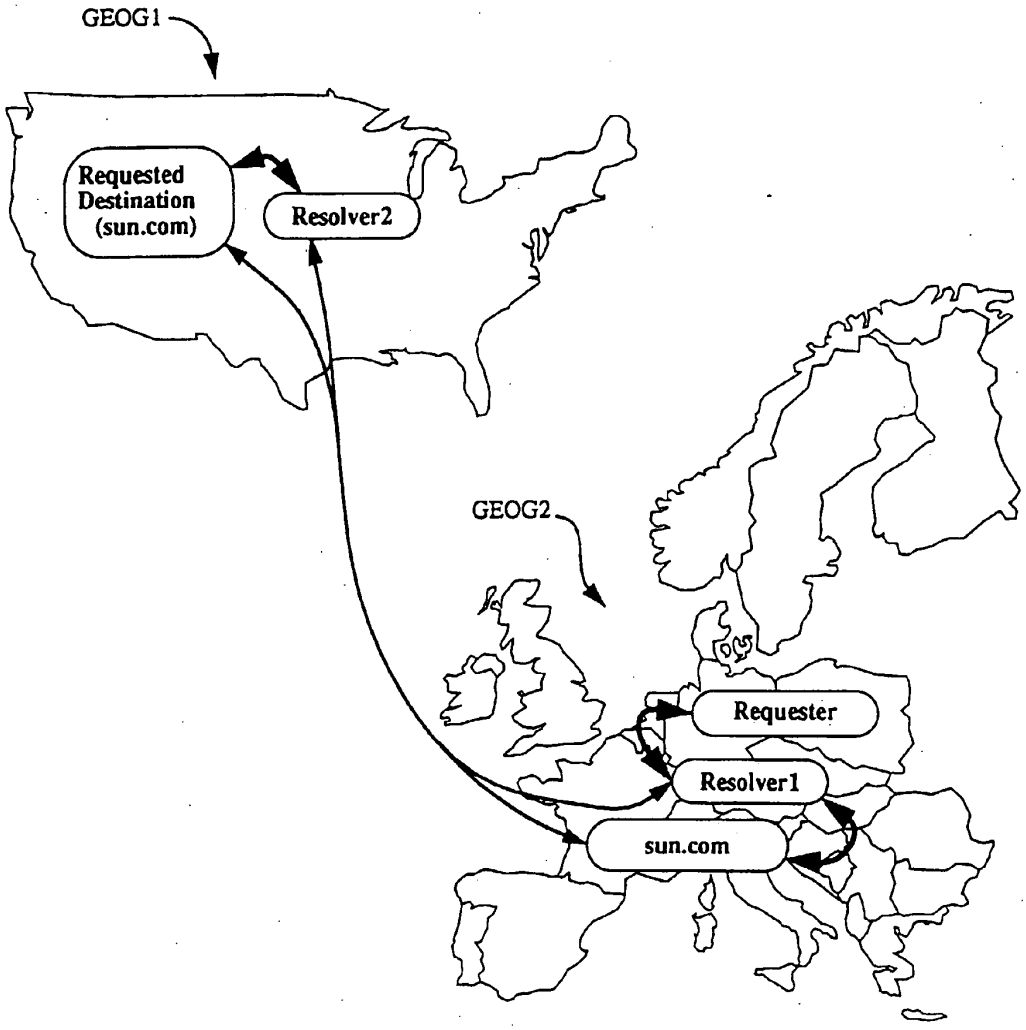

GEOG1
Requested Destination (sun.com)
Resolver2
GEOG2
Requester
Resolver1
sun.com

Figure 7

Zone 1 SENDER	Zone 2 RESOLVER SYSTEM	Zone 3 RECEIVER
Service Location Resolver1	Service Location Resolver2	Service Location Resolver3
Destination Lookup1	Destination Lookup2	

Citation 1

(19)日本国特許庁（ＪＰ）　(12)公開特許公報（A）　(11)特許出願公開番号

特開平10－27148

(43)公開日　平成10年(1998)１月27日

(51)Int.Cl.⁶		識別記号	庁内整理番号	FI		技術表示箇所
G06F	13/00	357		G06F	13/00	357Z
	15/16	470			15/16	470J
H04L	12/58		9744－5K	H04L	11/20	102D

審査請求　未請求　請求項の数11　OL　（全 11 頁）

(21)出願番号　特願平8－180722

(22)出願日　平成８年(1996)７月10日

(71)出願人　000005108
株式会社日立製作所
東京都千代田区神田駿河台四丁目６番地

(72)発明者　大野　修司
神奈川県川崎市幸区鹿島田890番地　株式会社日立製作所情報・通信開発本部内

(72)発明者　後藤　法宏
神奈川県横浜市戸塚区戸塚町5030番地　株式会社日立製作所ソフトウェア開発本部内

(72)発明者　小林　一可
神奈川県海老名市下今泉810番地　株式会社日立製作所オフィスシステム事業部内

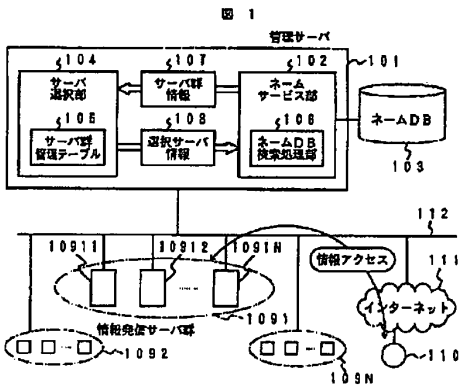
(74)代理人　弁理士　富田　和子

最終頁に続く

(54)【発明の名称】　インターネット用サーバシステム

(57)【要約】
【課題】端末利用ユーザに特別な操作を強いること無く、サーバ群内におけるアクセス分散と障害発生時等のバックアップが図れるインターネット用サーバシステムを提供することにある。
【解決手段】サーバ群１０９１に対して与えられたホスト名と、サーバ群１０９１を構成する複数台のサーバのそれぞれの通信用アドレスと、これらの通信用アドレスの中から目的の通信用アドレスを選択する際に用いられる選択情報とが格納されたサーバ群管理テーブル１０５と、インターネット１１１を介して、前記ホスト名に対応する通信用アドレスを要求されると、前記選択情報の内容に応じて、サーバ群管理テーブル１０５に格納されている通信用アドレスの何れかを選択するサーバ選択１０４と、サーバ選択部１０４が選択した通信用アドレスを通信用アドレスの要求元に送信するネームサービス部１０２とを備える。

図１



1

【特許請求の範囲】

【請求項1】インターネットを介して情報発信を行う複数台のサーバと、該複数台のサーバのそれぞれを管理する管理サーバとを備えたサーバシステムにおいて、

前記複数台のサーバの少なくとも一部は、同一情報の発信を行う複数台のサーバから成るサーバ群を構成し、

前記管理サーバは、

前記サーバ群に対して与えられたホスト名と、前記サーバ群を構成する前記複数台のサーバのそれぞれの通信用アドレスと、これらの通信用アドレスの中から目的の通信用アドレスを選択する際に用いられる選択情報とが格納された管理テーブルと、

前記インターネットを介して、前記ホスト名に対応する通信用アドレスを要求されると、前記選択情報の内容に応じて、前記管理テーブルに格納されている通信用アドレスの何れかを選択する選択手段と、

前記選択手段が選択した通信用アドレスを前記通信用アドレスの要求元に送信するネームサービス手段とを備えたことを特徴とするサーバシステム。

【請求項2】請求項1において、

前記選択情報には、前記通信用アドレスを選択するための選択順序が含まれており、

前記選択手段は、前記選択情報の前記選択順序にしたがって前記通信用アドレスを選択することを特徴とするサーバシステム。

【請求項3】請求項1において、

前記管理サーバは、各サーバの負荷状況を判断し、その判断結果を前記選択情報の一部として前記管理テーブルに格納する負荷監視手段をさらに備え、

前記選択手段は、前記通信用アドレスを選択する際に、前記判断結果を用いて、その時点で負荷が最も少ないサーバを選び、そのサーバの通信用アドレスを選択することを特徴とするサーバシステム。

【請求項4】請求項1において、

前記各サーバには、そのサーバの負荷状況を判断して、その結果を判断情報として前記管理サーバに送信する手段が設けられており、

前記管理サーバには、前記各サーバから送信されたそれぞれの判断情報を前記選択情報の一部として前記管理テーブルに格納する手段が設けられており、

前記選択手段は、前記通信用アドレスを選択する際に、前記判断情報を用いて、その時点で負荷が最も少ないサーバを選び、その通信用アドレスを選択することを特徴とするサーバシステム。

【請求項5】請求項1、2、3または4において、

前記サーバ群は、複数群構成されており、

前記管理テーブルには、前記各サーバ群に対して与えられたホスト名と、前記各サーバ群を構成する前記複数台のサーバのそれぞれの通信用アドレスと、これらの通信用アドレスの中から目的の通信用アドレスを選択する際に用いられる選択情報とが格納され、

前記選択手段は、前記インターネットを介して、前記ホスト名に対応する通信用アドレスを要求されると、前記選択情報の内容に応じて、該ホスト名で示されるサーバ群に含まれている各サーバの通信用アドレスの中から目的の通信用アドレスを選択することを備えたことを特徴とするサーバシステム。

【請求項6】インターネットを介して情報発信を行う複数台のサーバと、該複数台のサーバのそれぞれを管理する管理サーバとを備えたサーバシステムにおいて、

前記複数台のサーバの少なくとも一部は、同一情報の発信を行う複数台のサーバから成るサーバ群を構成し、

前記管理サーバは、

前記インターネット上にて前記サーバ群を一意に表すサーバ群用識別子と、前記インターネット上にて、前記サーバ群を構成する前記複数台のサーバのそれぞれを一意に表すサーバ用識別子と、これらのサーバ用識別子の中から目的のサーバ用識別子を選択する際に用いられる選択情報とが格納された管理テーブルと、

前記インターネットを介して、前記サーバ群用識別子を用いたアクセス要求を受付けると、前記選択情報の内容に応じて前記サーバ用識別子の何れかを選択する選択手段と、

前記選択手段が選択したサーバ用識別子を用いて前記サーバへのアクセスを行い、該アクセスで得られた情報を前記アクセス要求の要求元に送信すると共に、この情報を一時的に保持し、前記サーバ群用識別子を用いたアクセス要求を再度受付けた場合には、保持している当該情報を当該アクセス要求の要求元に送信する代理アクセス手段とを備えることを特徴とするサーバシステム。

【請求項7】請求項6において、

前記サーバ群用識別子は、前記サーバ群に与えられたホスト名であり、

前記複数のサーバ用識別子のそれぞれは、各サーバに与えられたホスト名であることを特徴とするサーバシステム。

【請求項8】請求項6または7において、

前記選択情報には、前記サーバ用識別子を選択するための選択順序が含まれており、

前記選択手段は、前記選択情報の前記選択順序にしたがって前記サーバ用識別子を選択することを特徴とするサーバシステム。

【請求項9】請求項6または7において、

前記管理サーバは、各サーバの負荷状況を判断し、その判断結果を前記選択情報の一部として前記管理テーブルに格納する負荷監視手段をさらに備え、

前記選択手段は、前記サーバ用識別子を選択する際に、前記判断結果を用いて、その時点で負荷が最も少ないサーバを選び、そのサーバのサーバ用識別子を選択することを特徴とするサーバシステム。

3

【請求項10】請求項6または7において、
前記各サーバには、そのサーバの負荷状況を判断して、その結果を判断情報として前記管理サーバに送信する手段が設けられており、
前記管理サーバには、前記各サーバから送信されたそれぞれの判断情報を前記選択情報の一部として前記管理テーブルに格納する手段が設けられており、
前記選択手段は、前記サーバ用識別子を選択する際に、前記判断情報を用いて、その時点で負荷が最も少ないサーバを選び、そのサーバ用識別子を選択することを特徴とするサーバシステム。
【請求項11】請求項6、7、8、9または10において、
前記サーバ群は、複数群構成されており、
前記管理テーブルには、前記各サーバ群に対して与えられたサーバ群用識別子と、前記各サーバ群を構成する前記複数台のサーバのそれぞれのサーバ用識別子と、これらのサーバ用識別子の中から目的のサーバ用識別子を選択する際に用いられる選択情報とが格納され、
前記選択手段は、前記インターネットを介して、前記サーバ群用識別子を用いたアクセス要求を受付けると、前記選択情報の内容に応じて、該サーバ群用識別子で示されるサーバ群に含まれている各サーバのサーバ識別子の中から目的のサーバ識別子を選択することを備えたことを特徴とするサーバシステム。
【発明の詳細な説明】
【0001】
【発明の属する技術分野】本発明は、インターネットを介して情報発信を行う複数台のサーバと、該複数台のサーバのそれぞれを管理する管理サーバとを備えたサーバシステムに関するものである。
【0002】
【従来の技術】インターネット上には 現在、ファイル転送のためのプロトコルであるＦＴＰ（Ｆｉｌｅ Ｔｒａｎｓｆｅｒ Ｐｒｏｔｏｃｏｌ）を利用したＦＴＰサーバや、テキストデータとグラフィックデータが混在したマルチメディア情報を扱うプロトコルであるＨＴＴＰ（Ｈｙｐｅｒ Ｔｅｘｔ Ｔｒａｎｓｆｅｒ Ｐｒｏｔｏｃｏｌ）を利用したＷＷＷ（Ｗｏｒｌｄ Ｗｉｄｅ Ｗｅｂ）サーバなどが情報発信ソースとして活躍している。
【0003】インターネットは、全世界の不特定多数のユーザがアクセス可能なネットワークで、現在急速な勢いで広まりつつあるが、一方では、人気の高いサーバにユーザからのアクセスが集中し、そのサーバに関して情報発信処理速度が著しく低下するという問題がある。また、目的のサーバが常に稼働しているとは限らず、故障やメインテナンス作業等のため、そのサーバにアクセスできないということもある。
【0004】このような状況に対処するため、インターネットに接続された情報発信サーバでは、従来より、以下に述べる２つの手法がよく用いられる。

4

【0005】第１の手法は、一つのサーバにアクセスが集中しないよう、同一の情報を複数台のサーバのそれぞれに持たせる方法である。
【0006】具体的には、図１０に示すような、ＦＴＰサーバのミラーサイトがこれに相当する。
【0007】同図において、オリジナルサイト５３のＦＴＰサーバＡ５４は、元情報を所有しており、ミラーサイト５１のＦＴＰサーバＢ５２は、オリジナルサイト５３のＦＴＰサーバＡ５４から該情報を定期的に引き出し、自らに格納する。これにより、端末１１０を利用して情報をアクセスしようとするユーザは、ＦＴＰサーバＡ５４が高負状態あるいは停止状態にあっても、ＦＴＰサーバＢ５２から必要な情報を手に入れることができるようになる。また、ＦＴＰサーバＡ５４から見れば、同サーバへのアクセスが分散される結果となり負荷が軽減される。
【0008】第２の手法は、情報発信サーバへのアクセス経路上に、情報発信サーバからユーザが引き出した情報を一時的に格納する（キャッシングする）サーバを配置し、その情報に対して、例えば、別のユーザがアクセスしようとした場合は、この情報を情報発信サーバからではなくキャッシングしているサーバから供給する方法である。
【0009】例えば、図１１に示すような、ＷＷＷサーバへのアクセス経路上に配置されたプロキシサーバがこれに相当する。
【0010】同図において、プロキシサーバ６２は、ＷＷＷサーバ６１の情報を得ようとする、端末１１０からのアクセス要求を受けると、端末１１０に代わって、ＷＷＷサーバ６１へアクセスする。その後、プロキシサーバ６２は、この代理アクセスで得た情報をキャッシュデータ６３として保持する共に、アクセス要求を発した端末１１０に該情報を返送する。その後、プロキシサーバ６２は、他の端末（図示省略）から該情報に対するアクセス要求を受付けると、該ＷＷＷサーバ６１にアクセスすることなく、自らが所有するキャッシュデータ６３を用いて返送処理を行う。これにより、ＷＷＷサーバ６１へのアクセス回数が軽減されることになる。
【0011】
【発明が解決しようとする課題】しかしながら、前述した第１の手法においては、各サーバの負荷軽減対策や、一つのサーバが停止した場合におけるバックアップ機能については実現されているものの、端末利用ユーザが必要な情報を得るためには、オリジナルサイトとミラーサイトのどちらかをユーザ側で選択しなければならない。このような選択作業は非常に煩雑であり、また、どのサーバが空いているのか等の判断をユーザに強いることになる。

【0012】また、前述した第2の手法においては、情報源となるサーバ（図11：WWWサーバ61）へのアクセス回数は軽減されるものの、該サーバの停止時においては、キャッシュ用サーバ（図11：プロキシサーバ62）が、これにアクセスすることができず、結果として、ユーザへの情報提供ができなくなる。

【0013】このような問題点に鑑み、本発明の目的は、端末利用ユーザに特別な操作を強いること無く、サーバ群内におけるアクセス分散と障害発生時等のバックアップが図れるインターネット用サーバシステムを提供することにある。

【0014】

【課題を解決するための手段】上記目的を達成するための本発明の一態様によれば、インターネットを介して情報発信を行う複数台のサーバと、該複数台のサーバのそれぞれを管理する管理サーバとを備えたサーバシステムにおいて、前記複数台のサーバの少なくとも一部は、同一情報の発信を行う複数台のサーバから成るサーバ群を構成し、前記管理サーバは、前記サーバ群に対して与えられたホスト名と、前記サーバ群を構成する前記複数台のサーバのそれぞれの通信用アドレスと、これらの通信用アドレスの中から目的の通信用アドレスを選択する際に用いられる選択情報とが格納された管理テーブルと、前記インターネットを介して、前記ホスト名に対応する通信用アドレスを要求されると、前記選択情報の内容に応じて、前記管理テーブルに格納されている通信用アドレスの何れかを選択する選択手段と、前記選択手段が選択した通信用アドレスを前記通信用アドレスの要求元に送信するネームサービス手段とを備えたことを特徴とするサーバシステムが提供される。

【0015】上記目的を達成するための本発明のその他の態様によれば、インターネットを介して情報発信を行う複数台のサーバと、該複数台のサーバのそれぞれを管理する管理サーバとを備えたサーバシステムにおいて、前記複数台のサーバの少なくとも一部は、同一情報の発信を行う複数台のサーバから成るサーバ群を構成し、前記管理サーバは、前記インターネット上にて前記サーバ群を一意に表すサーバ群用識別子と、前記インターネット上にて、前記サーバ群を構成する前記複数台のサーバのそれぞれを一意に表すサーバ用識別子と、これらのサーバ用識別子の中から目的のサーバ用識別子を選択する際に用いられる選択情報とが格納された管理テーブルと、前記インターネットを介して、前記サーバ群用識別子を用いたアクセス要求を受付けると、前記選択情報の内容に応じて前記サーバ用識別子の何れかを選択する選択手段と、前記選択手段が選択したサーバ用識別子を用いて前記サーバへのアクセスを行い、該アクセスで得られた情報を前記アクセス要求の要求元に送信すると共に、この情報を一時的に保持し、前記サーバ群用識別子を用いたアクセス要求を再度受付けた場合には、保持している当該情報を当該アクセス要求の要求元に送信する代理アクセス手段とを備えることを特徴とするサーバシステムが提供される。

【0016】

【発明の実施の形態】以下、本発明の一実施形態について説明する。

【0017】図1には、本実施形態のサーバシステムの全体構成が示されている。

【0018】このサーバシステムは、インターネット111を介して情報発信を行う複数の情報発信サーバ群（情報発信サーバ群1091、1092、…、109N）と、これらの情報発信サーバ群に対して行われる端末からのアクセスを管理する管理サーバ101とを有する。情報発信サーバ群1091、1092、…、109Nは、それぞれ、複数台のサーバから構成される。情報発信サーバ群1091は、同一情報を有する情報発信サーバ10911～1091Nから構成される。その他の情報発信サーバ群についても、これと同様に、そのサーバ群特有の情報をそれぞれの情報発信サーバが所有する。なお、特定の情報発信サーバ群へのアクセスが増加した場合は、同一情報を有するサーバを適宜追加すればよい。情報発信サーバ群1091、1092、…、109Nを管理するための管理情報は、ネームデータベース（ネームDB）103に格納されている。ネームDB103のデータ構造は、図2に示されている。

【0019】ネームDB103には、サーバネームフィールド21とサーバアドレスフィールド22の2つの格納フィールドが存在する。本実施形態では、あらかじめ、情報発信サーバ群1091、1092、…、109Nのそれぞれに対してホスト名を与えており、サーバネームフィールド21には、これらのホスト名が格納されている。このホスト名は、言うなれば、その情報発信サーバ群の各情報発信サーバを代表するホスト名であり、以下、このホストを便宜上、代表ホスト名と呼ぶこととする。

【0020】一方、本サーバシステムに設けられている全ての情報発信サーバは、それぞれ固有のサーバアドレスを有しており、サーバアドレスフィールド22には、これらのサーバアドレスが情報発信サーバ群毎にまとまった形で格納されている。サーバネームフィールド21の「－」は、その上の欄に格納されている内容と同じであることを意味している。情報発信サーバのアドレスは、インターネット111上で通信を行う際に使用する通信用アドレスで、具体的には、ＴＣＰ／ＩＰ通信を行うためのＩＰアドレスである。

【0021】なお、情報発信サーバ群1091、1092、…、109Nのそれぞれに与えられる代表ホスト名は、通常のホスト名、すなわちＤＮＳ（Ｄｏｍａｉｎ Ｎａｍｅ Ｓｙｓｔｅｍ）上で一義的に定まるホスト名と何らか変るものではない。

【0022】ここで、ドメイン（Domain）について簡単に説明する。

【0023】インターネットの構築がはじまった時点では、ネットワークに接続された個々のコンピュータを識別するために、そのホスト名とIPアドレスは、所定の組織で一元管理されていた。当初は、コンピュータの数が数百台のオーダであり、すべてのコンピュータのホスト名とIPアドレスを一つのファイルに登録することも可能であった。

【0024】しかしながら、ネットワークの拡大とともに、すぐにこの方法の限界が明らかになった。すなわち、次々に行われる新規登録や更新に莫大な時間を要すると共に、管理元の作業量が増大し、結果として収拾がつかなくなってしまった。

【0025】ドメインは、このような問題を解決するために提案された概念で、世界中に存在する数多くのコンピュータを一意に名前付けするために導入された。

【0026】インターネットのドメイン空間は、階層的に管理されており、DNSは、この管理体系にマッチするように設計されている。インターネットでは、各ドメインの情報をもつDNSサーバが分散配置され、ドメイン空間と同様に階層的に管理されている。本実施形態の管理サーバも、このDNSサーバの一種と考えることができる。

【0027】図2においては、情報発信サーバ群1091を示す代表ホスト名（serv bbb aaa）と、情報発信サーバ群1092を示す代表ホスト名（other bbb. aaa）が図示されている。サーバ群管理テーブル105には、これら以外の情報発信サーバ群の代表ホスト名についても全て登録されている。情報発信サーバ10911のサーバアドレス192. 128. 16. 1をはじめ、情報発信サーバ群1091を構成する各情報発信サーバのサーバアドレスは、情報発信サーバ群1091の代表ホスト名（serv. bbb aaa）に対応付けられた形で格納されている。情報発信サーバ群1092を構成する各情報発信サーバのサーバアドレスは、情報発信サーバ群1092の代表ホスト名（other. bbb. aaa）に対応付けられた形で格納されている。

【0028】管理サーバ101は、外部からのネームサービス要求を受けて、その応答処理を行うネームサービス部102と、ネームサービス部102にあってネームDB103から該当情報の検索・抽出を行うネームDB検索処理部106と、ネームサービス部102から、ネームDB検索処理部106によるネームDB103の検索結果であるサーバ群情報107を受け取り、サーバ群情報107が示すサーバ群からサーバ1台を選択し、その結果を選択サーバ情報108としてネームサービス部102に通知するサーバ選択部104と、サーバを選択する際にサーバ選択部104が参照するサーバ群管理テーブル105とを備える。

【0029】以上の構成を有する管理サーバ101は、情報発信サーバ群1091～109Nと同一LAN（Local Area Network）112上に存在しており、このLAN112がインターネット111に接続している。インターネット111には無数の端末が接続しており、各ユーザは、それぞれの端末を用いて本サーバシステムの各情報発信サーバ群をアクセスすることが可能である。なお、図1には、代表的に端末110が図示されている。

【0030】つぎに、端末110から情報発信サーバ群1091をアクセスする場合を例にとって、本サーバシステムの動作について説明する。ここでは、情報発信サーバ群1091に付与されている代表ホスト名をユーザが予め知っているものとして話しを進める。

【0031】ユーザは、まず、情報発信サーバ群1091の代表ホスト名（serv bbb. aaa）を端末110に入力する。端末110は、この入力操作を受けて、代表ホスト名に対応するサーバアドレスを管理サーバ101に問い合わせる。具体的には、この際、代表ホスト名（serv bbb. aaa）を示す情報が管理サーバ101に送られる。ここまでは、従来のDNSサーバにネームサービスを要求する場合と同じである。

【0032】管理サーバ101のネームサービス部102は、端末110からのアドレス問い合わせ、すなわち、代表ホスト名（serv. bbb aaa）を示す情報を受け取ると、ネームDB検索処理部106を起動する。ネームDB検索処理部106は、図5に示すネームDB検索処理を実行する。

【0033】同図に示すように、ネームDB検索処理部106は、まず、端末110から送られた代表ホスト名をキーにネームDB103を検索する（ステップ51（S51））。これにより、キーにした代表ホスト名に対応する全てのサーバアドレスが抽出される。ここでは、代表ホスト名（serv bbb. aaa）に対応するサーバアドレス（192 128. 16 1、192 128. 16. 4、 ）が抽出される。次に、ネームDB検索処理部106は、見つかったサーバアドレスの全てを先程の代表ホスト名と共にサーバ群情報107としてサーバ選択部104に通知する（S52）。サーバ選択部104は、サーバ群管理テーブル105を参照し、その内容に応じて、情報提供を行わせるに最適なサーバを一つ選択する。この選択処理の詳細は後述することとする。選ばれたサーバのホスト名とサーバアドレスは、選択サーバ情報108として、ネームサービス部102に送られる。ネームサービス部102は、このサーバアドレスを検索結果として処理し、その後、アドレス問い合わせ要求の応答処理として、このサーバアドレスを端末110に送信する（S53）。

【0034】管理サーバ101からサーバアドレスを通

9

知された端末１１０は、これ以後、そのサーバに対して直接、情報アクセスを行えるようになる。

【００３５】このように本実施形態によれば、或る情報発サーバ群がアクセスされる度に、その情報発信サーバ群の中から、情報発信を行うに最適なサーバが選択されるので、結果として、端末からのアクセスがその情報発信サーバ群内で分散されることになる。

【００３６】なお、特に説明しなかったが、端末１１０では、該端末１１０から送出された代表ホスト名が管理サーバ１０１に直接、あるいは、ここでは図示しないＤＮＳを介して到達するよう、あらかじめ環境設定が為されている。これについては、ＤＮＳ上で通常行われていることなので詳細は省略する。

【００３７】また、上記実施形態では、同一ＬＡＮに情報発信サーバ群と管理サーバを接続しているが、これは、情報発信サーバ群が管理サーバと同一ＬＡＮ上に存在しなくてはならないことを意味するものではない。例えば、情報発信サーバ群の各サーバはインターネット上の任意の場所に接続されていてもよい。

【００３８】つぎに、本発明のその他の実施形態について説明する。

【００３９】図３には、本実施形態のサーバシステムの全体構成が示されている。

【００４０】本サーバシステムの管理サーバ３０１は、インターネットでさまざまなマルチメディア情報の発信を行うＷＷＷ（Ｗｏｒｌｄ　Ｗｉｄｅ　Ｗｅｂ）サーバ等へのアクセスにおいて、端末からＷＷＷサーバへのアクセスを直接受け付け、端末の代わりにＷＷＷサーバに対する情報アクセス、及び、そのアクセス情報のキャッシングを行う、いわゆるプロキシサーバの役割を果たすものである。

【００４１】図中、１０９１～１０９Ｎは、同一情報を有するサーバ複数台からなる情報発信サーバ群である。ここでは、情報発信サーバ群１０９１が、複数台のＷＷＷサーバから構成され、情報発信サーバ群１０９２が、複数台のＦＩＰサーバから構成されている。

【００４２】管理サーバ３０１において、３０２は、端末からの情報アクセスに応じて、情報発信サーバへの代理アクセス、アクセス情報の端末への返送、及び、そのアクセス情報の一時保持（キャッシュ処理）を行う代理アクセス部、３０６は、代理アクセス部３０２にあってサーバリスト３０３から該当情報の検索・抽出を行うサーバリスト検索処理部、３０４は、代理アクセス部３０２から、サーバリスト検索処理部３０６によるサーバリスト３０３の検索結果であるサーバ群情報３０７を受け取り、その情報が示すサーバ群の中からサーバ１台を選択し、その結果を代理アクセス部３０２に選択サーバ情報３０８として通知するサーバ選択部、３０５は、サーバ選択部３０４にあってサーバの選択処理の際に参照されるサーバ群管理テーブルである。

10

【００４３】サーバリスト３０３のデータ構造は、図４に示されている。

【００４４】サーバリスト３０３には、図４に示すように、代表サーバネームフィールド４１とサーバ群内サーバネームフィールド４２の２つの格納フィールドが存在する。本実施形態でも、前述と同様に、情報発信サーバ群１０９１、１０９２、…、１０９Ｎのそれぞれに対してホスト名（代表ホスト名）を与えており、サーバネームフィールド４１には、これらの代表ホスト名が格納されている。代表サーバネームフィールド４１の「－」は、その上の欄に格納されている内容と同じであることを意味している。また、サーバ群内サーバネームフィールド４２には、代表サーバネームフィールド４１に格納されている代表ホスト名で示される情報発信サーバ群を構成する各サーバのホスト名が格納されている。

【００４５】図４には、情報発信サーバ群１０９１を示す代表ホスト名（ｗｗｗ　ｂｂｂ．ａａａ）と、情報発信サーバ群１０９２を示す代表ホスト名（ｆｔｐ　ｂｂｂ．ａａａ）が図示されている。情報発信サーバ１０９１１のホスト名（ｓｅｒｖ１．ｗｗｗ－ｃ．ｂｂｂ．ａａａ）をはじめ、情報発信サーバ群１０９１を構成する各情報発信サーバのホスト名は、情報発信サーバ群１０９１の代表ホスト名（ｗｗｗ　ｂｂｂ．ａａａ）に対応付けられた形で格納されている。情報発信サーバ群１０９２を構成する各情報発信サーバのホスト名は、情報発信サーバ群１０９２の代表ホスト名（ｆｔｐ　ｂｂｂ　ａａａ）に対応付けられた形で格納されている。

【００４６】情報発信サーバ群１０９１～１０９Ｎとアクセス管理サーバ３０１は、それぞれ同一ＬＡＮ（Ｌｏｃａｌ　Ａｒｅａ　Ｎｅｔｗｏｒｋ）１１２に接続している。同ＬＡＮ１１２は、インターネット１１１に接続している。インターネット１１１には、情報発信サーバ群１０９１～１０９Ｎが発信している情報を入手する際に使用するユーザの端末１１０が接続している。

【００４７】つぎに、情報発信サーバ群１０９１の各サーバが発信している情報を端末１１０で入手する場合を例にとり、本サーバシステムの動作を説明する。ここでは、情報発信サーバ群１０９１に付与されている代表ホスト名をユーザが予め知っているものとして話しを進める。

【００４８】ユーザは、まず、情報発信サーバ群１０９１の代表ホスト名（ｗｗｗ．ｂｂｂ　ａａａ）を端末１１０に入力する。端末１１０は、この入力操作を受けて、代表ホスト名を管理サーバ３０１に送り、情報発信サーバ群１０９１へのアクセスを要求する。ここまでは、従来のプロキシサーバにアクセスを要求する場合と同じである。

【００４９】管理サーバ３０１の代理アクセス部３０２は、端末１１０から代表ホスト名（ｗｗｗ．ｂｂｂ．ａａａ）を受け取ると、サーバリスト検索処理部３０６を

起動する。サーバリスト検索処理部306は、図6に示すサーバリスト検索処理を実行する。

【0050】サーバリスト検索処理部306は、まず端末110から送られた代表ホスト名をキーにサーバリスト303を検索する（ステップ61（S61））。これにより、キーにした代表ホスト名に対応する全てのホスト名が抽出される。ここでは、代表ホスト名（serv．bbb．aaa）に対応するホスト名（serv1．www－c bbb aaa、serv4．www－c．bbb aaa、…）が抽出される。次に、サーバリスト検索処理部306は、見つかったホスト名全てを先程の代表ホスト名と共にサーバ群情報307としてサーバ選択部304に通知する（S62）。サーバ選択部304は、サーバ群管理テーブル305を参照し、その内容に応じて、情報提供を行わせるに最適なサーバを一つ選択する。この選択処理の詳細は後述することとする。選ばれたサーバのホスト名は、前記代表ホスト名と共に選択サーバ情報308として、代理アクセス部302に送られる。代理アクセス部302は、このホスト名を用いて、そのサーバにアクセスする。代理アクセス部302は、この代理アクセスによって得られた情報を該管理サーバ301内のメモリにキャッシュデータとして格納し、さらに、このキャッシュデータを端末110に返送する。次回、代表ホスト名（www bbb aaa）によるアクセス要求を受け付けた場合には、このキャッシュデータが当該アクセス要求元に送信される。

【0051】以上の動作により、端末110は、管理サーバ301を介して間接的に情報発信サーバにアクセスする。

【0052】このように本実施形態によれば、管理サーバ301が行う代理アクセスが、その時点で最適なサーバに対して行われるため、結果として、そのサーバ群内のアクセスが分散される。

【0053】以上、本発明のサーバシステムの2つの実施形態について、その全体構成と動作を中心に説明したが、続いて、管理サーバの備えるサーバ選択部の詳細を説明する。

【0054】ここでは、便宜上、図1に示したサーバシステムを例にとって話しを進める。

【0055】まず、図7に示したサーバ群管理テーブル105を用いた場合のサーバ選択部104の動作について説明する。

【0056】ネームサービス102からサーバ選択部104に送られるサーバ群情報107には、ネームDB検索処理（図5参照）でキーとした代表ホスト名と、該代表ホスト名をキーとして検索された検索結果である複数のサーバアドレスが含まれているが、このうち、図7のサーバネームフィールド1051には、前記代表ホスト名が格納される。また、サーバアドレスフィールド1052には、検索結果である前述の複数のサーバアドレスが格納される。

【0057】また、サーバ選択部104が選択したサーバのホスト名は、選択サーバ情報108としてネームサービス102に送られるが、選択サーバ情報108で通知したサーバアドレスについては、アクセスフラグフィールド1053に「1」（1：通知）が設定される。なお、その他のサーバアドレスについては、「0」（0：未通知）を設定される。

【0058】そして、サーバ選択部104は、サーバ群情報107を受け取ると、サーバ群情報107の示す代表ホスト名がサーバネームフィールド1051に存在するか否かを判断し、該代表ホスト名がサーバネームフィールド1051に無ければ、この代表ホスト名と共に、これに対応付けられている全てのサーバアドレス（サーバ群情報107に含まれている全てのサーバアドレス）をサーバ群管理テーブル105に登録する。このとき、一つのサーバアドレスについては、そのアクセスフラグフィールド1053に"1"を設定する。他のアクセスフラグフィールド1053には、"0"を設定する。同図においては、サーバアドレス192．128．16．1のアクセスフラグフィールド1053に"1"を設定し、これ以外は"0"を設定している。

【0059】一方、すでに同一代表ホスト名がサーバネームフィールド1051に格納されている場合は、つぎの作業を行う。

【0060】例えば、サーバ群管理テーブル105の内容が図7の状態に既に設定されている場合、サーバアドレス192．128　16．4のアクセスフラグフィールド1053に"1"を設定すると共に、サーバアドレス192　128　16　1のアクセスフラグフィールド1053に"0"を設定する。なお、最後尾に位置する、サーバアドレス192．128　16．12のアクセスフラグフィールド1053に"1"が設定されていた場合は、先頭行に位置する、サーバアドレス192　128．16　1のアクセスフラグフィールド1053に"1"を設定すると共に、サーバアドレス192．128　16．12のアクセスフラグフィールド1053に"0"を設定する。

【0061】以上の処理を行った場合、設定値である"1"は、アクセスフラグフィールド1053を順番に移動していくかのように見える。

【0062】つぎに、サーバ選択部104は、アクセスフラグフィールド1053内に"1'を設定した行のサーバアドレスフィールド1052内に格納されているサーバアドレスを選択サーバ情報108としてネームサービス部102に通知する。

【0063】このように情報発信を行わせるサーバを次々に選択していけば、端末からのアクセスは、その情報発信サーバ群内の各サーバに対してほぼ均等に分散されるようになる。

【0064】つぎに、図8に示したサーバ群管理テーブル105を用いた場合のサーバ選択部104の動作について説明する。

【0065】ここでは、情報発信サーバ群を構成する各サーバの負荷に応じて、アクセスするサーバを選択する。

【0066】図8において、同じ構成要素については、同じ番号を付して説明を省略する。

【0067】このサーバ群管理テーブル105は、先程のアクセスフラグフィールド1053に代えて、負荷レベルフィールド1054を備えている点に特徴がある。負荷レベルフィールド1054には、各サーバの負荷状況を示す情報が格納される。

【0068】また、選択部104では、新たなにサーバ負荷監視部113が設けられている。サーバ負荷監視部113は、サーバ群管理テーブル105に登録されているサーバ（すなわち、アドレスフィールド1052にサーバアドレスが格納されているサーバ）の負荷状況を定期的に監視し、負荷レベルフィールド1054にその負荷レベルを設定する。負荷レベルは予め数段階（ここでは、負荷レベルA、B、Cの3段階）に分けられている。具体的には、サーバ負荷監視部113は、サーバアドレスフィールド1052に格納されているサーバアドレスのサーバに対して定期的に監視コマンドを送信し、その応答時間を集計し、その結果を、負荷レベルA／B／Cの何れかに分けて負荷レベルフィールド1054に設定する。ここで、上記監視コマンドに対して何も応答しないサーバについては、その負荷レベルフィールド1053に対して何も設定しない。サーバが何も応答しない場合、そのサーバは、保守作業中、あるいは、障害の発生による停止状態にあることが予想される。

【0069】そして、サーバ選択部104は、サーバ群情報107を受けた時点で、サーバ群管理テーブル105に登録されているサーバの中から最も負荷の軽いものを選択し、このサーバのアドレス等を選択サーバ情報108としてネームサービス部102に通知する。この際、負荷レベルフィールド1053に何も設定されていないサーバについては、あらかじめ選択枝から除外しておくこととする。

【0070】このような処理を行えば、端末からのアクセスは、目的の情報発信サーバ群の中から、その時点で最も負荷の軽いサーバに対して行われるようになる。また、稼働していないサーバへのアクセスも禁止されることになる。

【0071】なお、各サーバの負荷に応じてアクセスするサーバを選択する場合、サーバシステムは図9に示すように構成しても構わない。

【0072】同図において、各情報発信サーバには、そのサーバの負荷をモニターし、その負荷状況を管理サーバ101に報告する負荷モニタ115が設けられている。管理サーバ101は、各サーバからの負荷状況報告を収集し、その結果を負荷レベルとしてサーバ群管理テーブル105に格納する負荷状況収集部114を備える。ここで、負荷モニタ部115から報告のないサーバについては負荷レベルを設定しない。

【0073】そして、サーバ選択部104は、サーバ群情報107を受け付けると、サーバ群管理テーブル105に登録されているサーバの中から負荷レベル情報が設定されているものを探し出し、その中から最も負荷の軽いものを選択し、このサーバのサーバアドレスを選択サーバ情報108としてネームサービス部102に通知する。

【0074】このように各サーバの負荷状況を、そのサーバ自身にモニターさせれば、より正確な負荷情報を得ることができるようになる。

【0075】以上、管理サーバの備えるサーバ選択部の各種実施形態について説明したが、これらは、図3に示したサーバシステムにも適用可能である。

【0076】

【発明の効果】本発明によれば、障害発生時などのバックアップ対策として、また、ユーザからのアクセスを分散させる目的で、同一の情報を有する複数台のサーバを設けた場合であっても、情報発信を行わせるに最適なサーバが自動的に選択されるため、ユーザは、複数台のサーバの存在を意識することなく、必要な情報を手に入れることができるようになる。

【0077】また、本発明によれば、情報発信を行わせるに最適なサーバは、各サーバの負荷状況に応じて、また、稼働していないサーバは対象外として選択されるため、ユーザは、常に最良の環境でアクセスを行うことができるようになる。

【0078】その上、アクセス負荷が増加しても、本発明を提供しておくことにより、同一情報を有するサーバを追加するだけで、アクセス負荷に応じた性能補強が可能となる。

【図面の簡単な説明】

【図1】本発明に係るサーバシステムの一実施形態の全体構成図。

【図2】図1に示した管理サーバが備えるネームDBの構造を示した説明図。

【図3】本発明に係るサーバシステムのその他の実施形態の全体構成図。

【図4】図3に示した管理サーバが備えるサーバリストの構造を示した説明図。

【図5】図1に示した管理サーバが行うネームDB検索処理の一例を示したフローチャート。

【図6】図3に示した管理サーバが行うサーバリスト検索処理の一例を示したフローチャート。

【図7】図1に示した管理サーバの選択部の一例（その1）を示した説明図。

15

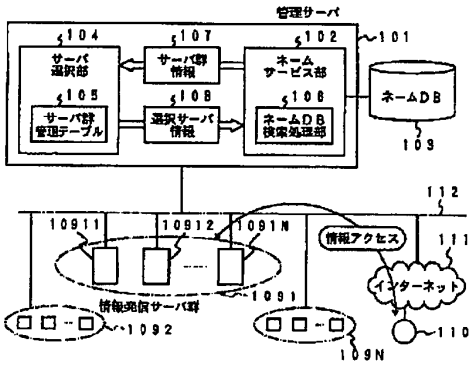
【図8】図1に示した管理サーバの選択部の一例（その2）を示した説明図。
【図9】本発明に係るサーバシステムのさらに別の実施形態の全体構成図。
【図10】従来のサーバシステム（ＦＴＰミラー）の一例を示した説明図。
【図11】従来のサーバシステム（プロキシアクセス）一例を示した説明図。
【符号の説明】
101、301…管理サーバ、
102　ネームサービス部、
103…ネームＤＢ、
104、304…サーバ選択部、
105、305　サーバ群管理テーブル、
106…ネームＤＢ検索処理部、

16

＊107、307…サーバ群選択情報、
108、308　選択サーバ情報、
1091～109N…情報発信サーバ群、
10911～1091N…情報発信サーバ、
110　端末、
111　インターネット、
112…LAN（Local Area Network）、
302…代理アクセス部、
303　サーバリスト、
306　サーバリスト検索処理部、
113…サーバ負荷監視部、
114　負荷状況収集部、
115…負荷モニタ部

【図1】

図 1



【図2】

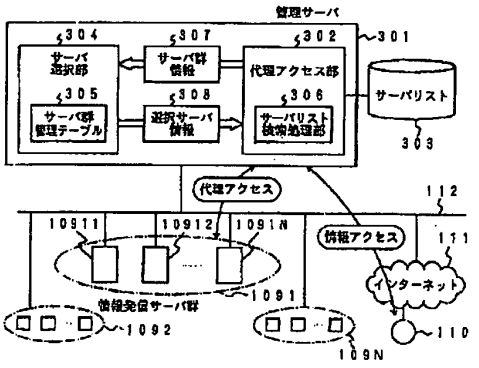
図 2

21	22
サーバネーム	サーバアドレス
serv.bbb.aaa	192.128.16.1
−	192.128.16.4
−	⋮
−	192.128.16.12
other.bbb.aaa	192.128.64.18
⋮	⋮

103

【図3】

図 3



【図4】

図 4

41	42
代表サーバネーム	サーバ群内サーバネーム
www.bbb.aaa	serv1.www-c.bbb.aaa
−	serv4.www-c.bbb.aaa
−	⋮
ftp.bbb.aaa	serv2.ftp-c.bbb.aaa
−	serv1.ftp-c.bbb.aaa
−	⋮

303

【図7】

図 7

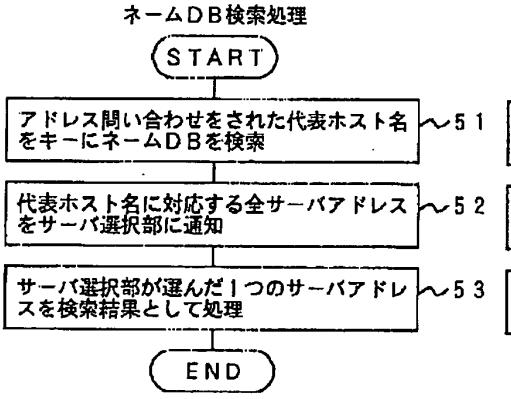
104

1051	1052	105 サーバ群管理テーブル
サーバネーム	サーバアドレス	アクセスフラグ
serv.bbb.aaa	192.128.16.1	1
−	192.128.16.4	0
−	⋮	⋮
−	192.128.16.12	0

1059

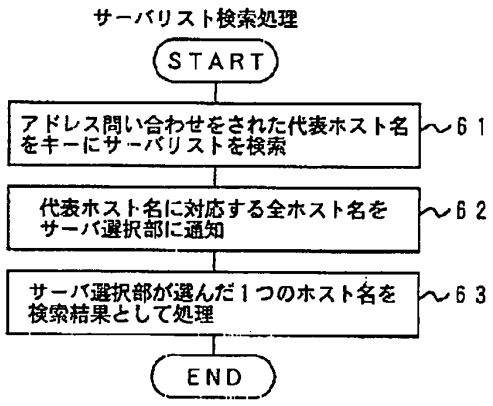
【図5】

図 5



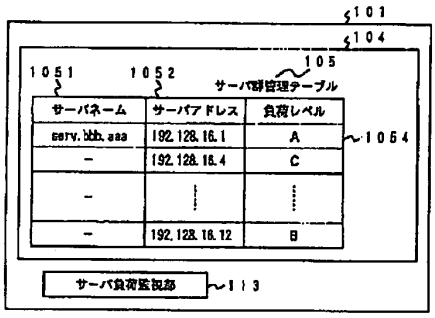
【図6】

図 6



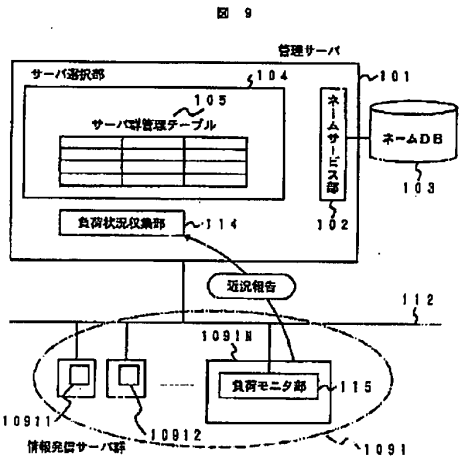
【図8】

図 8



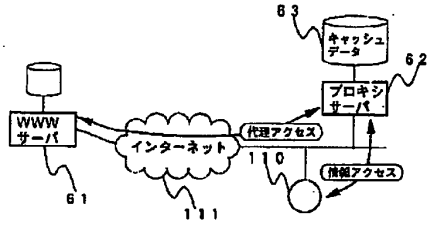
【図9】

図 9



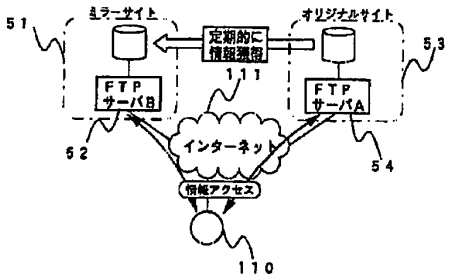
【図11】

図11



【図10】

図10



フロントページの続き

(72)発明者　綿貫　達哉
神奈川県川崎市幸区鹿島田890番地　株式会社日立製作所情報・通信開発本部内

(19)日本国特許庁（ＪＰ）　(12) **公開特許公報（A）**　(11)特許出願公開番号

特開平7－66829

(43)公開日　平成7年(1995)3月10日

(51)Int.Cl.[6]	識別記号	庁内整理番号	ＦＩ	技術表示箇所
Ｈ04Ｌ 12/54				
Ｇ06Ｆ 13/00	351 Ｇ	7368－5B		
Ｈ04Ｌ 12/58				
		8732－5K	Ｈ04Ｌ 11/20　101 Ｂ	

審査請求　未請求　請求項の数5　ＯＬ　（全 8 頁）

(21)出願番号　特願平5－162529

(22)出願日　平成5年(1993)6月30日

(71)出願人　000001258
川崎製鉄株式会社
兵庫県神戸市中央区北本町通1丁目1番28号

(72)発明者　斎藤　洋
東京都千代田区内幸町2丁目2番3号　川崎製鉄株式会社東京本社内

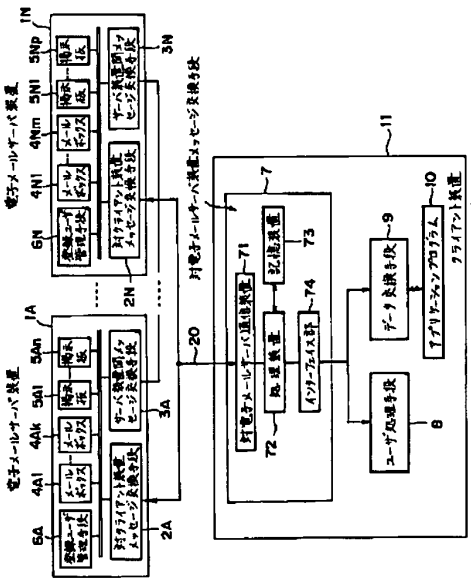
(74)代理人　弁理士　谷　義一　(外1名)

(54)【発明の名称】　電子メール多重化システムおよび該システムにおける通信制御方法

(57)【要約】

【構成】　多重化された電子メールサーバ装置（1A－1N）がそれぞれ有する、メールボックス（または、電子掲示板）の内容を、相互に複写する機能を、各電子メールサーバ装置にもたせ、クライアント装置（11）内の対電子メールサーバ装置メッセージ交換手段（7）が、稼働中の電子メールサーバ装置の1台を自動的に選択して、両者を接続するようにした。また、電子メールサーバ装置からの情報に基づいて、最も負荷の少ない電子メールサーバ装置をクライアント装置が選択する。

【効果】　複数の電子メールサーバ装置のうち、少なくとも1台が稼働状態であれば、サービスが継続される。多重化された電子メールサーバ装置間の負荷分散が可能である。



1

【特許請求の範囲】

【請求項1】　多重化された電子メールサーバ装置とクライアント装置とがネットワークを介して接続された電子メールシステムにおいて、

前記電子メールサーバ装置は、

メールボックスまたは電子掲示板等のメッセージ蓄積手段と、

登録されたユーザの管理をする登録ユーザ管理手段と、

前記クライアント装置とのメッセージの送受信を行う対クライアント装置メッセージ交換手段と、

前記電子メールサーバ装置間でのメッセージの交換を行うサーバ装置間メッセージ交換手段とを備え、

前記クライアント装置は、

使用者またはアプリケーションプログラムとのインターフェイスをとる手段と、

前記電子メールサーバ装置とのメッセージの送受信を行う対電子メールサーバ装置メッセージ交換手段とを備え、

前記サーバ装置間メッセージ交換手段は、複数の電子メールサーバ装置間で、前記メッセージ蓄積手段の内容を複写し、

前記対電子メールサーバ装置メッセージ交換手段は、前記電子メールサーバ装置からの情報に基づいて、前記電子メールサーバ装置の中から、稼働中の装置を通信相手として選択することによって、電子メールシステムの無停止を実現することを特徴とする電子メール多重化システム。

【請求項2】　請求項1において、前記対電子メールサーバ装置メッセージ交換手段は、前記電子メールサーバ装置からの情報に基づいて、負荷の少ない電子メールサーバ装置を選択することを特徴とする電子メール多重化システム。

【請求項3】　多重化された電子メールサーバ装置とクライアント装置とがネットワークを介して接続された電子メールシステムにおいて、

前記電子メールサーバ装置間でメッセージの交換を行い、複数の電子メールサーバ装置間で、前記メッセージを複写する第1の過程と、

前記クライアント装置が、前記電子メールサーバ装置からの情報に基づいて、前記電子メールサーバ装置の中から、稼働中の装置を通信相手として選択する第2の過程とを具備することを特徴とする電子メール多重化システムにおける通信制御方法。

【請求項4】　請求項3において、前記第2の過程は、

前記クライアント装置が前記電子メールサーバ装置に対して自局のＩＤを一斉送信する過程と、

該送信に応じた複数の電子メールサーバ装置の中から一つの電子メールサーバ装置を選択する過程と、

選択した電子メールサーバ装置に接続要求する過程とを具備することを特徴とする電子メール多重化システムにおける通信制御方法。

2

【請求項5】　請求項4において、前記クライアント装置が、通信相手の電子メールサーバ装置の故障を検出した場合には、前記一斉送信に応じた複数の電子メールサーバ装置の中から稼働中の他の装置を選択して、接続要求を送信することを特徴とする電子メール多重化システムにおける通信制御方法。

【発明の詳細な説明】

【０００１】

【産業上の利用分野】本発明は、多重化された電子メールサーバ装置を備えた電子メール多重化システムおよび該システムにおける通信制御方法に関する。

【０００２】

【従来の技術】従来の電子メールシステムとしては、特開平４－２７３６３６号公報に記載された分散掲示板方式が知られている。これは、電子メールシステムにおいて、複数のサーバ装置に掲示板のコピーを配置することにより、掲示物の読み出し要求が１つのサーバ装置に集中することを防止し、応答性を向上させようとしたものである。

【０００３】この従来の電子メールシステムにおいては、電子メールサーバ装置とクライアント装置とが、ネットワークを介して、固定的に接続されていた。すなわち、特定の利用者が、電子メールシステムから必要なデータを得る場合は、特定の１台の電子メールサーバ装置からそのデータを得ていた。なお、この種のデータを利用するものとしては、一般に、利用者とアプリケーションプログラムとがあるが、以下の説明においては、これらを総称して、「サービス受益対象」と呼ぶこととする。

【０００４】

【発明が解決しようとする課題】従来の電子メールシステムにおいては、特定のサービス受益対象と電子メールサーバ装置とが固定的に接続されているため、自分が接続された電子メールサーバ装置が停止した場合、サービス受益対象は、この電子メールシステムを利用することができなかった。このような事態を避けるためには、電子メールシステムの無停止サービスを実現しなければならない。具体的には、電子メールサーバ装置にフォールトトレラント型のコンピュータを適用したり、複数の電子メールサーバ装置を高度に結合し、多重化する方法がとられていた。これらの装置を実現するには、高度な制御技術と、多大な費用が要求された。

【０００５】さらに、同一のネットワーク上に複数の電子メールサーバ装置が稼働していても、特定の電子メールサーバ装置に、アクセスが集中してしまう可能性があった。

【０００６】本発明は、無停止で稼働する安価な電子メール多重化システムを提供することを目的とする。また、多重化された電子メールサーバ装置間の負荷分散を

可能にしたシステムを提供することを目的とする。

【0007】

【課題を解決するための手段】上記課題を解決するために、本発明は、多重化された電子メールサーバ装置に、それぞれが有するメールボックス（または、電子掲示板）の内容を、相互に複写する機能をもたせるとともに、クライアント装置が、稼働中の電子メールサーバ装置の1台を自動的に選択して接続するようにした。すなわち、本発明では、電子メールサーバ装置の切り替えを、クライアント装置が行うようにした。

【0008】また、電子メールサーバ装置からの情報に基づいて、最も負荷の少ない電子メールサーバ装置をクライアント装置が選択することによって、複数の電子メールサーバ装置間の負荷のバランスをとっている。

【0009】

【作用】本発明によれば、複数の電子メールサーバ装置のうち、少なくとも1台が稼働状態であれば、クライアント装置内に配置されている対電子メールサーバ装置メッセージ交換手段が、稼働中の電子メールサーバ装置の中の1台を自動的に選択して接続を行う。

【0010】さらに、クライアント装置が接続する際に、電子メールサーバ装置側からの情報に基づいて、負荷の最も少ない電子メールサーバ装置に接続することによって、多重化された電子メールサーバ装置間の負荷分散を可能にした。

【0011】

【実施例】次に、本発明の実施例を詳細に説明する。

【0012】図1は、本発明の一実施例を示す図である。

【0013】図1において、電子メールサーバ装置1A－1Nは、メッセージを蓄積して、要求に応じてメッセージを要求元に転送する装置である。また、クライアント装置11は、サービス受益対象が、電子メールサーバ装置1A－1Nにサービスを要求するための装置である。

【0014】電子メールサーバ装置1A－1Nは、それぞれ、対クライアント装置メッセージ交換手段2A－2N、サーバ装置間メッセージ交換手段3A－3N、メールボックス4A1－4Ak－4N1－4Nm、掲示板5A1－5An－5N1－5Np、および登録ユーザ管理手段6A－6Nを備えている。

【0015】ここで、各メールボックスおよび掲示板は、メッセージ蓄積手段として機能するもので、サーバ装置間メッセージ交換手段3A－3Nを介して、相互に接続されている。登録ユーザ管理手段6A－6Nは、電子メールサーバ装置1A－1Nに接続を要求してきたサービス受益対象の認証を行うもので、内部にユーザIDとパスワードとを格納している。

【0016】クライアント装置11は、対電子メールサーバ装置メッセージ交換手段7と、ユーザ処理手段8と、データ交換手段9と、アプリケーションプログラム10とを備えている。ここで、対電子メールサーバ装置メッセージ交換手段7は、各電子メールサーバ装置との通信を制御する対電子メールサーバ通信装置71と、処理装置72と、プログラムおよび各種データを格納する記憶装置73と、インターフェイス部74とを有している。対電子メールサーバ通信装置71は、処理装置72からの要求によって、図2および図3に示す通信制御を行う。また、インターフェイス部74は、処理装置72とユーザ処理装置8およびデータ交換手段9とのインターフェイスをとる。

【0017】ユーザ処理手段8は、利用者からのメッセージや指示を受け取り、電子メールサーバ装置1A－1Nに対して、メッセージを転送したり、メッセージの転送を要求したり、また転送されてきたメッセージを利用者に供給する。電子メールサーバ装置との、これらのやり取りは、クライアント装置内の対電子メールサーバ装置メッセージ交換手段7、ネットワーク20、電子メールサーバ装置内の対クライアント装置メッセージ交換手段2A－2Nを介して行われる。

【0018】アプリケーションプログラム10は、クライアント装置11内に、必要に応じて配置されたもので、データ交換手段9および対電子メールサーバ装置メッセージ交換手段7を通して、電子メールサーバ装置1A－1Nと接続され、メッセージ交換を行う。

【0019】次に、電子メールサーバ装置1（1A－1N）の多重化について説明する。

【0020】ある電子メールサーバ装置内のメールボックス4および掲示板5は、サーバ装置間メッセージ交換手段3を通して、他の電子メールサーバ装置内のメールボックス4および掲示板5にコピーされる。このコピーは、ネットワークの構成や、メッセージの転送頻度などに応じて決められる。具体的には、夜間や休日など、装置が比較的すいている間に、一括コピーする方法や、メッセージがいずれかの電子メールサーバ装置に到着した時点で、ただちに、他の電子メールサーバ装置のメールボックスや掲示板にコピーする方法などがある。なお、コピーの数は、必ずしも同一である必要はなく、メッセージの重要性に応じて、変えてもよい。

【0021】次に、図2および図3を参照して、本実施例の動作を説明する。

【0022】図2は、接続時の手順を示すものであり、サービス受益対象（ユーザまたはアプリケーションプログラム）、対電子メールサーバ装置メッセージ交換手段7、および電子メールサーバ装置の3つに分けて、動作を示している。

【0023】サービス受益対象は、電子メールサーバ装置1からのサービスを受けるために、自身のIDおよびパスワードを、クライアント装置11に指示する（ステップS1－S4）。クライアント装置11のユーザ処理

5

手段8またはデータ交換手段9は、対電子メールサーバ装置メッセージ交換手段7に対して、ＩＤとパスワードとを引き渡した上で、電子メールサーバ装置1への接続を依頼する。このとき、対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象のＩＤとパスワードとを内部的に記憶しておく。

【００２４】次に、対電子メールサーバ装置メッセージ交換手段7は、接続可能なすべての電子メールサーバ装置1に対して、サービス受益対象のＩＤと特定の指示とを、一斉送信する（ステップＳ５）。ここで、特定の指示とは、上記ＩＤをもつサービス受益対象にサービス提供可能な電子メールサーバ装置は、その電子メールサーバ装置の識別コードを返送するように、との指示である。

【００２５】一斉送信を受信した電子メールサーバ装置1Ａ－1Ｎは、受信したサービス受益対象のＩＤと、自分が管理する登録ユーザ管理手段6Ａ－6Ｎ内のデータとを照合し、一致するものがあれば、自分の識別コードを、対クライアント装置メッセージ交換手段2Ａ－2Ｎを経由して、問い合わせ元の対電子メールサーバ装置メッセージ交換手段7に返す（ステップＳ６－Ｓ９）。こうして、このサービス受益対象にサービスを提供できる電子メールサーバ装置の識別コードが、対電子メールサーバ装置メッセージ交換手段7に受信される（ステップＳ１０）。

【００２６】対電子メールサーバ装置メッセージ交換手段7は、サービスを提供できる電子メールサーバ装置のリストを作成し（ステップＳ１１）、接続すべき電子メールサーバ装置を選択する（ステップＳ１３）。この場合、各電子メールサーバ装置に対して、現在の負荷状況などの問い合わせをして（ステップＳ１２）、電子メールサーバ装置を決定してもよい。一方、サービス提供可能な電子メールサーバ装置が存在しない場合、対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象に対して、サービス提供ができない旨の通知を行う（ステップＳ１４－Ｓ１５）。サービス受益対象は、この通知を受信して、接続が失敗であったことを認識し、処理を終了する（ステップＳ１６）。

【００２７】利用可能な電子メールサーバ装置が複数存在する場合、対電子メールサーバ装置メッセージ交換手段7は、あらかじめ定められた手順で、一台の電子メールサーバ装置を選択する。具体的には、次のような方法が考えられる。

【００２８】（1）　サービス可能なすべての電子メールサーバ装置に対して、処理負荷に関する問い合わせを行い（ステップＳ１２）、その時点で負荷が最も少ない電子メールサーバ装置を選択する。

【００２９】（2）　対電子メールサーバ装置メッセージ交換手段7に、最初に識別コードを返した電子メールサーバ装置を選択する。

6

【００３０】電子メールサーバ装置の選択が完了すると、対電子メールサーバ装置メッセージ交換手段7は、選択した電子メールサーバ装置に対して、ログイン要求、ユーザＩＤ、およびパスワードを送信する（ステップＳ１７）。電子メールサーバ装置は、これらを受信し（ステップＳ１８）、接続が可能か否かを確認のうえ（ステップＳ１９）、接続を実行する（ステップＳ２０）。接続完了後、あるいは接続が不可能な場合、電子メールサーバ装置は、対電子メールサーバ装置メッセージ交換手段7に対して、結果を通知する（ステップＳ２１）。

【００３１】対電子メールサーバ装置メッセージ交換手段7は、接続結果を受信し（ステップＳ２２）、接続が成功したときには、サービス開始をサービス受益対象に通知する（ステップＳ２３－Ｓ２４）。サービス受益対象は、このサービス開始通知を受信し、次の段階（サービス要求）に入る。一方、接続が失敗に終った場合は、ステップＳ１３に戻って、他の電子メールサーバ装置を選択し、この電子メールサーバ装置に接続要求を行う。指定した電子メールサーバ装置との接続に再度失敗した場合は、この電子メールサーバ装置を、サービス可能な電子メールサーバ装置のリストから外して、他の電子メールサーバ装置にもう一度接続要求を行う。

【００３２】このようにして、電子メールサーバ装置との接続に成功するか、あるいは、サービス可能な電子メールサーバ装置がなくなるまで、ステップＳ１３－Ｓ２３の手順を繰り返す。

【００３３】図3は、正常時のサービス要求の動作と、電子メールサーバ装置に障害が発生した場合の動作とを示すフローチャートである。

【００３４】電子メールサーバ装置との接続に成功すると、サービス受益対象は、対電子メールサーバ装置メッセージ交換手段7に対して、サービス要求を送信する（ステップＳ３１）。このサービス要求を受信すると、対電子メールサーバ装置メッセージ交換手段7は、電子メールサーバ装置に対して、サービス要求を送信する（ステップＳ３２）。サービス要求を受信した電子メールサーバ装置は、サービスを実行し、サービス内容を対電子メールサーバ装置メッセージ交換手段7に返送する（ステップＳ３３－Ｓ３５）。

【００３５】サービス内容を受信した対電子メールサーバ装置メッセージ交換手段7は、サービス内容をサービス受益対象に転送する（ステップＳ３６－Ｓ３７）。サービス受益対象はこのサービス内容を受領して処理を終了する（ステップＳ３８）。電子メールサーバ装置との接続に成功し、電子メールサーバ装置からのサービスの提供が可能になった場合、電子メールサーバ装置と対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象が電子メールサービスが不要となるまで、その接続を継続する。もし、何らかの原因で、接続中の電子

7

メールサーバ装置がサービスを提供できなくなった場合は、次の障害発生時のフローにしたがって、サービスの再要求が行われる。

【００３６】対電子メールサーバ装置メッセージ交換手段７が、電子メールサーバ装置に対してサービス要求を送信した後、あらかじめ定めた時間、サービス内容が返送されない場合、対電子メールサーバ装置メッセージ交換手段７は、サービス内容の受信に失敗したものと判断し（ステップＳ４１）、電子メールサーバ装置に対してサービスの再要求を行う。すなわち、ステップＳ１３－Ｓ２３において、電子メールサーバ装置を選択し、選択した電子メールサーバ装置に対して、サービス再要求を送信する（ステップＳ３２）。なお、この手順は、図２の接続時の手順と同様である。この手順は、サービス受益対象に何ら通知せずに行われる。すなわち、サービスを提供する電子メールサーバ装置が障害を起こすたびに、対電子メールサーバ装置メッセージ交換手段７は、他のサービス可能な電子メールサーバ装置とクライアント装置とを順次接続し、サービスの要求およびサービス内容の転送を実行する（ステップＳ３２－Ｓ３８）。こうして、無停止の電子メールシステムを実現することができる。

【００３７】ただし、サービス可能な電子メールサーバ装置が無くなった場合には、対電子メールサーバ装置メッセージ交換手段７は、サービス受益対象に対して、サービス提供不能通知を送信し（ステップＳ１５）、サービス受益対象は、電子メールサーバ装置が障害である旨の報告を行う（ステップＳ１６ａ）。

【００３８】

8

【発明の効果】以上説明したように、この発明は、電子メールサーバ装置の切り替えをクライアント装置側で行うことによって、安価な電子メールシステムを実現することができる。また、処理負荷の最も軽い電子メールサーバ装置を選択する方式を使用することによって、負荷分散を実現することができる。

【図面の簡単な説明】

【図１】この発明による電子メール多重化システムの実施例の構成を示すブロック図である。

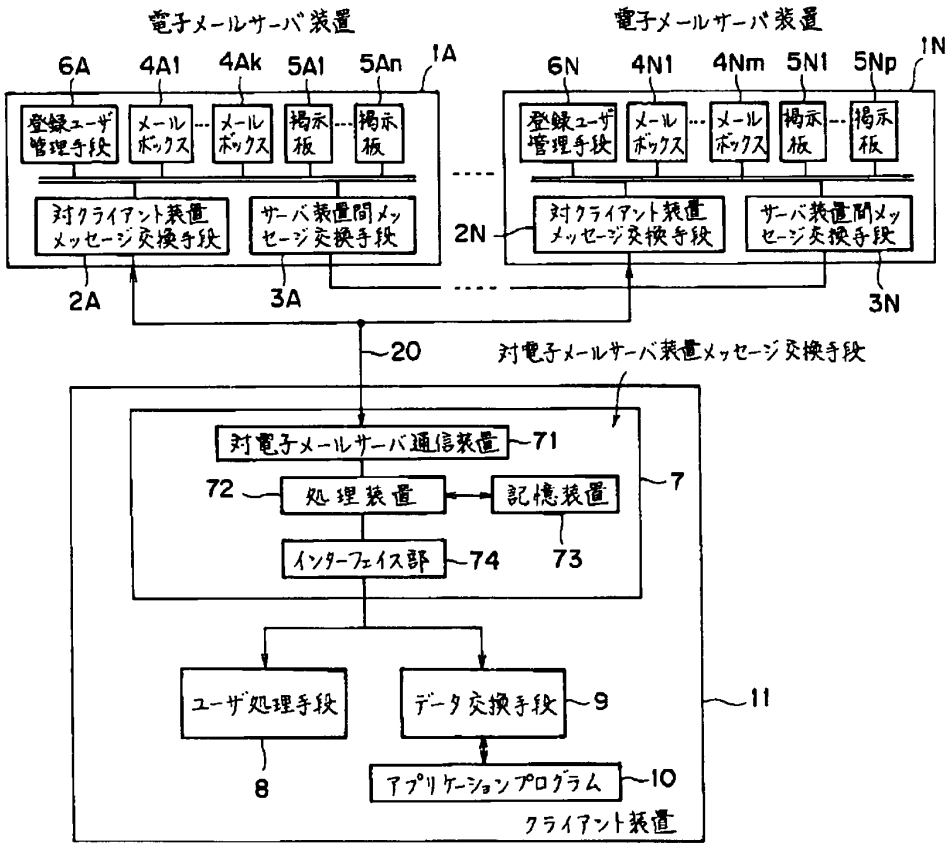
【図２】上記実施例の接続時の動作を示すフローチャートである。

【図３】上記実施例のサービス要求時の動作を、正常時と、障害発生時とに分けて示すフローチャートである。

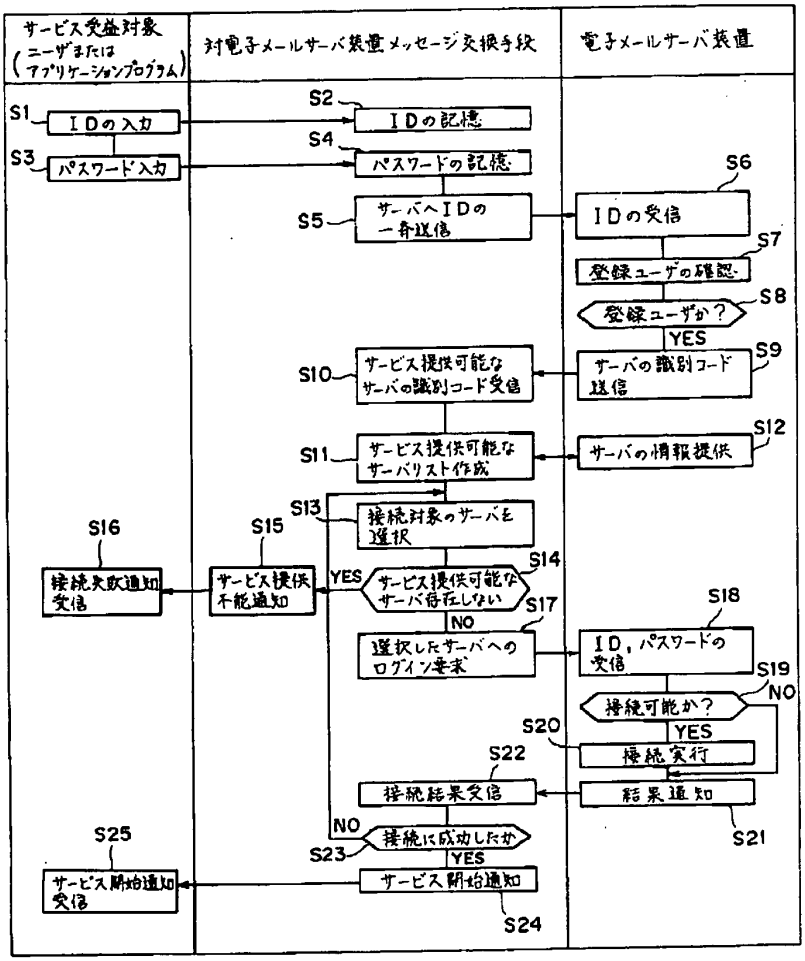
【符号の説明】

１Ａ，１Ｎ　電子メールサーバ装置
２Ａ，２Ｎ　対クライアント装置メッセージ交換手段
３Ａ，３Ｎ　サーバ装置間メッセージ交換手段
４Ａ１－４Ａｋ，４Ｎ１－４Ｎｍ　メールボックス
５Ａ１－５Ａｎ，５Ｎ１－５Ｎｐ　掲示板
６Ａ－６Ｎ　登録ユーザ管理手段
７　対電子メールサーバ装置メッセージ交換手段
８　ユーザ処理手段
９　データ交換手段
１０　アプリケーションプログラム
２０　ネットワーク
７１　対電子メールサーバ通信装置
７２　処理装置
７３　記憶装置
７４　インターフェイス部

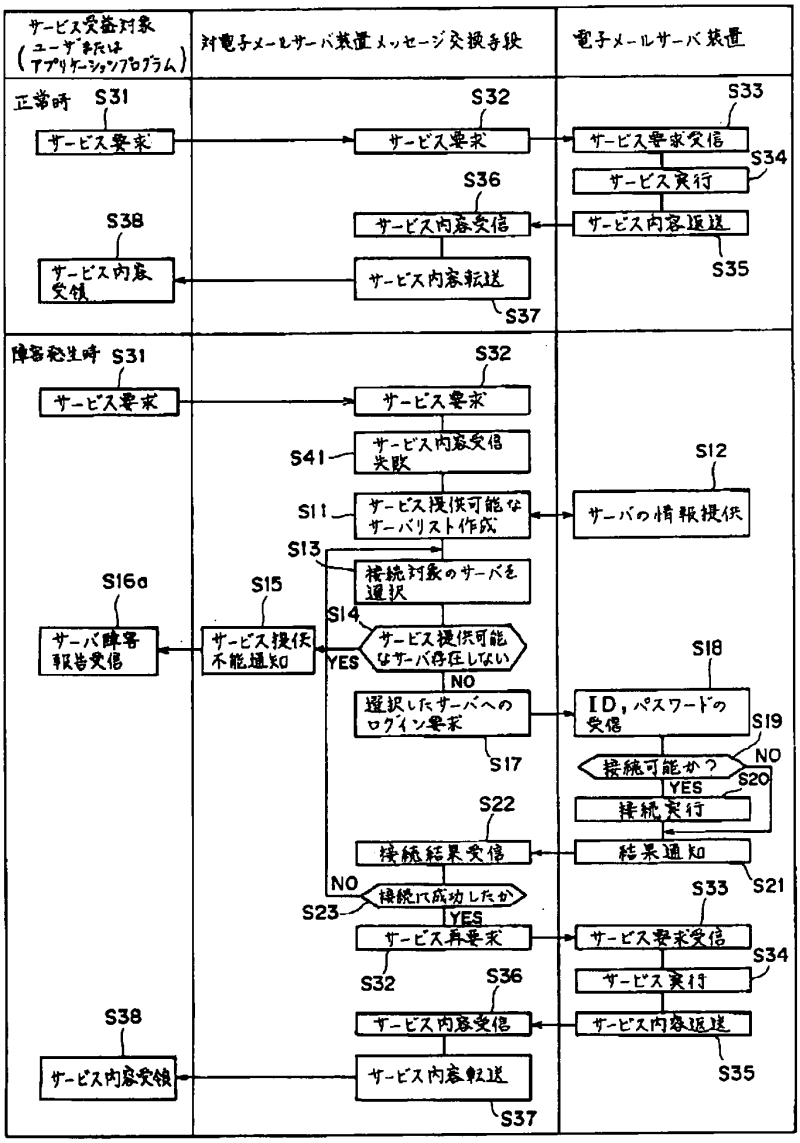
【図1】


電子メールサーバ装置
電子メールサーバ装置
6A
4A1
4Ak
5A1
5An
1A
6N
4N1
4Nm
5N1
5Np
1N
登録ユーザ管理手段
メールボックス
メールボックス
掲示板
掲示板
対クライアント装置メッセージ交換手段
サーバ装置間メッセージ交換手段
2N
2A
3A
3N
20
対電子メールサーバ装置メッセージ交換手段
対電子メールサーバ通信装置
71
72
処理装置
記憶装置
7
インターフェイス部
74
73
ユーザ処理手段
データ交換手段
9
11
8
アプリケーションプログラム
10
クライアント装置

【図2】


サービス受益対象
(ユーザまたは
アプリケーションプログラム)
対電子メールサーバ装置メッセージ交換手段
電子メールサーバ装置
S1
IDの入力
S2
IDの記憶
S3
パスワード入力
S4
パスワードの記憶
S5
サーバへIDの
一斉送信
S6
IDの受信
S7
登録ユーザの確認
S8
登録ユーザか？
YES
S9
サーバの識別コード
送信
S10
サービス提供可能な
サーバの識別コード受信
S11
サービス提供可能な
サーバリスト作成
S12
サーバの情報提供
S13
接続対象のサーバを
選択
S14
サービス提供可能な
サーバ存在しない
YES
S15
サービス提供
不能通知
S16
接続失敗通知
受信
NO
S17
選択したサーバへの
ログイン要求
S18
ID，パスワードの
受信
S19
接続可能か？
NO
YES
S20
接続実行
S21
結果通知
S22
接続結果受信
NO
S23
接続に成功したか
YES
S24
サービス開始通知
S25
サービス開始通知
受信

【図3】


サービス受益対象（ユーザ系またはアプリケーションプログラム）
対電子メールサーバ装置メッセージ交換手段
電子メールサーバ装置
正常時
S31
サービス要求
S32
サービス要求
S33
サービス要求受信
S34
サービス実行
S36
サービス内容受信
サービス内容返送
S35
S38
サービス内容受領
サービス内容転送
S37
障害発生時
S31
サービス要求
S32
サービス要求
S41
サービス内容受信失敗
S12
S11
サービス提供可能なサーバリスト作成
サーバの情報提供
S13
接続対象のサーバを選択
S16a
S15
S14
サーバ障害報告受信
サービス提供不能通知
YES
サービス提供可能なサーバ存在しない
S18
NO
選択したサーバへのログイン要求
ID，パスワードの受信
S19
S17
接続可能か？
NO
YES
S20
接続実行
S22
接続結果受信
結果通知
S21
NO
S23
接続に成功したか
S33
YES
サービス再要求
サービス要求受信
S34
S32
S36
サービス実行
サービス内容受信
サービス内容返送
S38
S35
サービス内容受領
サービス内容転送
S37

HTTP/1.1 200 OK
Date: Tue, 09 Apr 2002 12:04:12 GMT
Server: Apache/1.3.20 (Unix)
Last-Modified: Wed, 01 Oct 1997 09:23:00 GMT
ETag: "3ddb76-3a9c-34321674"
Accept-Ranges: bytes
Content-Length: 15004
Connection: close
Content-Type: text/plain

INTERNET-DRAFT
<draft-vinod-carp-v1-01.txt>

J. Cohen, N. Phadnis
Netscape Communications
Vinod Valloppillil
Microsoft Corporation
Keith W. Ross
University of Pennsylvania
29 September 1997
Expires March 1998

Cache Array Routing Protocol v1.1

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

This draft documents proposed changes to the original version of the draft with this same name, to be called 'the Cache Array Routing Protocol (CARP) v1.1' for dividing URL-space among an array of loosely coupled proxy servers.

The modified sections are 3.1 and 3.2. Aside from those two sections, the drafts remains the same.

An overview of the changes are presented in section 0.

An HTTP client agent (either a proxy server or a client browser) which implements CARP v1.1 can allocate and intelligently route requests for the correct URLs to any member of the Proxy Array. Due to the resulting sorting of requests through these proxies, duplication of cache contents is eliminated and global cache hit rates are improved.

INTERNET-DRAFT Cache Array Routing Protocol v1.1 30 June 1997

Table of Contents

0. Proposed Change Overview..... +
1. Overview..... 2
2. Proxy Array Membership Table..... 3
   2.1 Global Information..... 3
   2.2 Member Information..... 4
3. Routing Function..... 5
   3.1 Hash Function..... 5
   3.2 Hash Combination..... 6
   3.3 Load Factor..... 6
   3.4 Route Selection..... 7
   3.5 Member Failure Routing..... 7
4. Client-Side Implementation..... 7
5. Versioning..... 7
6. Security Considerations..... 7
7. Open Issues..... 7
8. Acknowledgements..... 7
9. References..... 8
10. Author's Information..... 8

0. Proposed Change Overview

Change in hash functions

The following changes are to make hash functions faster, better (resulting in more uniform distribution), and easier to evaluate (especially in JavaScript on client side).

1. Convert a URL to lower case to evaluate hash value to avoid parsing the URL. This makes hashing faster at client as well as server side.

2. Use shift instead of rotate operator. It is four times faster and gives more uniform distribution based on our experiments.

3. Additive multiplication is replaced with a simple multiplication. Additive effect can be achieved simply by adding 1 to the multiplier constant if required. Since the multiplier constant is anyway a set of random bits, adding 1 to it has no effect toward improving distribution.

4. The second rotate operation in evaluation of member proxy hash values as well as resultant hash value is eliminated. It is a relatively expensive operation and does not help improve distribution after the multiplication.

1. Overview

The Cache Array Routing Protocol describes a distributed caching protocol based on

1) a known membership list of loosely coupled proxies and
2) a hash function for dividing URL space among those proxies

The Proxy Array Membership Table is defined as a plain ASCII text file retrieved from an Array Configuration URL. This document does NOT describe how this table is constructed, merely the format of the fields used by agents implementing.

The hash function plus routing algorithm defined in this document take member proxies described in the Proxy Array Membership Table and make an on-the-fly determination as to which Proxy Array member should be the proper receptacle for a cached version of a resource keyed by URL.

Downstream agents may then access the cached resource by forwarding the proxied HTTP request [5] for that resource to the appropriate member of the Proxy Array.

INTERNET-DRAFT Cache Array Routing Protocol v1.1 30 June 1997

2. Proxy Array Membership Table

The Proxy Array Membership Table is a plain-text ASCII file which can be published from a URL.

The format of the table is:

Proxy Array Information/<Version number>

```
ArrayEnabled:  <0 | 1>
ConfigID:  <opaque string>
ArrayName:  <opaque string>
ListTTL:  <minutes until next check>

<name> <IP addr> <listening port> <table URL> <agent str>
<statetime> <status UP | DOWN> <load factor> <cache size>
```

2.1 Global Information

These are fields that describe the array itself and are not specific to any one member of an array

Global information is terminated in the Proxy Array Membership Table by a CR/LF/CR/LF.

2.1.1 Version number

The version number for implementations of this specification is 1.0

2.1.2 ArrayEnabled

This field allows proxies to advertise their implementation of CARP i v1 even if they are not members of a Proxy Array.

2.1.3 ConfigID

ConfigID is an opaque number no larger than 32bits similar to an ETag in HTTP 1.1. It is used to track the current state of an Array table and may be used to match multiple yet independently published copies of the Proxy Array Membership Table.

2.1.4 ArrayName

ArrayName is an opaque string which is used to provide a convenient administrative name for a given array.

2.1.5 ListTTL

ListTTL is the number of seconds for which an HTTP client entity should consider the current table image valid. After ListTTL has expired, that client should retrieve a new copy of the Proxy Array Membership Table.

2.2 Member Information

The following fields are published per member in an array and are separated by single spaces. The end of an array member's record is terminated by a CR/LF.

2.2.1 Name

The name of the proxy server. Typically this is the fully qualified DNS name. Downstream HTTP agents should use resolution of this name to determine how to connect to this proxy.

2.2.2 IP Addr

The IP address that other proxy servers within this array should use to connect to this proxy server. This is necessary for proxy servers which may be hosted on multi-hommed servers where requests are only accepted by one of the interfaces.

2.2.3 Listening Port

The TCP port number this proxy is expecting requests on.

2.2.4 Table URL

A URL which may be maintained by this proxy server on which a copy of the array membership table can be found.

2.2.5 Agent String

An opaque string identifying the vendor / version of the proxy Server in the Array Membership Table.

2.2.7 Statetime

How long a Proxy Server has been in its current state and has been a member of this table. This is useful for dynamic generation of the Array Membership Table where the host generating the table has knowledge of the proxy's operational status.

2.2.8 Status

Status provides a simple text string indicating whether a member proxy is currently able to handle requests (UP) or refused a connection when last contacted (DOWN).

2.2.9 Load Factor

Load Factor is a relative amount of the total load for an array that should be handled by any given member of the array.

2.2.10 Cache Size

Cache size is an informational field that indicates the size of the cache held by a particular member of an array.

3. Routing Function

Once an agent has a Proxy Array Membership Table. It uses a mathematical hash function to determine which of the members of the array should be the receptacle of a particular URL request.

This routing function involves constructing n "scores" using a hash of the request URL plus a hash of each of the n proxies in the Proxy Array Membership Table.

Both the URL and the proxy names are hashed in order to minimize the disruption of target routes if a member of the target array can't be contacted.

Hashes of the URL and proxy name are constructed using the algorithm described in 3.1 and combined using the algorithm described in 3.2.

3.1. Hash Function

The hash function outputs a 32 bit unsigned integers based on a zero-terminated ASCII input string. Machine names as well as URLs should be converted to lower case for hash evaluation. This makes sense because machine names are case insensitive.

Because irreversibility and strong cryptographic features are unnecessary for this application, a very simple and fast hash function based on the bitwise left shift operator is used.

```
For (each char in URL):
      URL_Hash += (URL_Hash << 9) + char ;
```

Member proxy hashes are computed in a similar manner:

```
For (each char in MemberProxyName):
      MemberProxy_Hash +=    (MemberProxy_Hash << 9) + char ;
```

Becaues member names are often similar to each other, their hash values are further spread across hash space via a multiplication:

```
MemberProxy_Hash = MemberProxy_Hash * 0x62531965 ;
```


3.2. Hash Combination

Hashes are combined by first exclusive or-ing (XOR) the URL hash by the machine name and then multiplying by a constant.

All final and intermediate values are 32 bit unsigned integers.

```
Combined_Hash = (URL_hash ^ MemberProxy_Hash) ;
Combined_Hash = Combined_Hash * 0x62531965 ;
```

3.3. Load Factor

Support for array members with differing HTTP processing & caching capacity is achieved by multiplying each of the combined hash values by a Load Factor Multiplier.

The Load Factor Multiplier for an individual member is calculated by taking each member's relative Load Factor and applying the following formula:

For each proxy server 1,...,K, the Load Factor Multiplier, X_k, is calculated iteratively as follows:

```
X_1 = arbitrary positive constant

X_k = ([K-k+1] * [P_k - P_{k-1}])/(X_1 * X_2 * ... * X_{k-1})
X_k += (X_{k-1}^{K-k+1})
X_k = X_k^{1/(K-k+1)}
```

where:

```
X_k = Load Factor Multiplier for proxy k
K = number of proxies in an array
P_k = relative percent of the load that proxy k should handle
```

This is then combined with the previously computed hashes as

Resultant_value = Combined_Hash * X_k

3.4. Route Selection

The "score" for a particular combination of URL plus proxy is its resultant value. Once the agent determines of the scores of the K proxies, it routes the URL query to the proxy with the hightest score.

3.5. Member Failure Routing

If a proxy can not contact the designated member of a proxy array in order to forward an HTTP request, that proxy should route the request to the second highest scoring proxy in the target array.

4. Client-side implementation

CARP can be implemented on client-side HTTP browsers via the use of the Proxy AutoConfig file described in [1] and [2].

5. Versioning

If a downstream proxy receives an Array Membership Table with a greater version # than that proxy is able to parse, it should fall back to simple proxy request routing to any administrator defined upstream proxy server.

6. Security Considerations

This draft does not discuss relevant security considerations.

7. Open Issues

8. Acknowledgements

The author would like to thank Brian Smith, Kip Compton, and Kerry Schwartz for their assistance in preparing this document.

Most of the architecture & design of CARP stem from work conducted by Brian Smith at Microsoft Corp.

9. References

[1] Luotonen, Ari., "Navigator Proxy Auto-Config File Format", Netscape Corporation, http://home.netscape.com/eng/mozilla/2.0/relnotes/demo/proxy-live.html, March 1996.

[2] Microsoft Corporation., "Automatic Proxy Configuration", http://www.microsoft.com/ie/ieak/autosys.htm, March 21, 1997.

[3] Wessels, Duane., "Internet Cache Protocol Version 2", http://ds.

internic.net/internet-drafts/draft-wessels-icp-v2-00.txt, March 21, 1997.

[4] Sharp Corporation., "Super Proxy Script", http://naragw.sharp.co.jp/sps/, August 9, 1996.

[5] Fielding, R., et. al, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2068, UC Irvine, January 1997.

[6] Valloppillil & Cohen, "Hierarchical HTTP Routing Protocol", http://ircache.nlanr.net/Cache/ICP/draft-vinod-icp-traffic-dist-00.txt, April 21, 1997.

10. Author Information

Josh Cohen <josh@netscape.com>
Nilkanth Phadnis <phadnis@netscape.com>
Netscape Communications
501 E. Middlefield Rd
Mountain View, Ca 94043

Vinod Valloppillil
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone: 1.206.703.3460
Email: VinodV@Microsoft.Com

Keith W. Ross
University of Pennsylvania
Department of Systems Engineering
Philadelphia, PA 19104

Phone: 1.215.898.6069
Email: Ross@UPenn.Edu

Expires December 1997

English Translation of Citation 1:

(prepared by the Japanese Patent Office)

Japanese Patent Publication No. 10(1998)-27148

Date of Publication: January 27, 1998

Japanese Patent Application No. 08(1996)-180722

Date of Application: July 10, 1996

Inventor: Ono, Shuji

Applicant: Hitachi Ltd.

Title of Invention: Server System for Internet

What follows is a MACHINE-TRANSLAION of Citation 1 that has been prepared by the Japanese Patent Office.

[Detailed Description of the Invention]

[0001]

[Field of the Invention] This invention relates to the server system equipped with two or more sets of the servers which perform information dispatch through the Internet, and the management server which manages each of the server of these two or more bases.

[0002]

[Description of the Prior Art] On the Internet, the FTP server using FTP (File Transfer Protocol) which is a protocol for current and a file transfer, the WWW (World Wide Web) server using HTTP (Hyper Text Transfer Protocol) which is a protocol treating the multimedia

information to which text data and graphical data were intermingled, etc. are playing an active part as the information dispatch source.

[0003] On the other hand, access from a user concentrates on a popular server, and the Internet has the problem that information dispatch processing speed falls remarkably about the server, although many and unspecified global users are accessible networks and are spreading with vigor with current [rapid]. Moreover, it may not restrict that the target server is always working, but it may be said for failure, a maintenance activity, etc. that the server cannot be accessed.

[0004] In order to cope with such a situation, in the information dispatch server connected to the Internet, two technique described below is used well conventionally.

[0005] The 1st technique is an approach of giving each of two or more sets of servers the same information, as access does not concentrate on one server.

[0006] Specifically, the mirror site of a FTP server as shown in drawing 10 is equivalent to this.

[0007] In this drawing, FTP server A54 of the original copy site 53 owns former information, and FTP server B-52 of a mirror site 51 pulls out this information periodically from FTP server A54 of the original copy site 53, and it stores it in oneself. Thereby, even if the user who is going to access information using a terminal 110 has FTP server A54 in a high negative condition or a idle state, he can get required information from FTP server B-52. Moreover, if

it sees from FTP server A54, a result by which access to this server is distributed will be brought, and a load will be mitigated.

[0008] The 2nd technique is the approach of supplying this information from the server which is carrying out the cash advance from the information dispatch server, when the server (a cash advance is carried out) which stores temporarily the information which the user pulled out from the information dispatch server tends to be arranged and another user tries to access on the access path to an information dispatch server as opposed to that information.

[0009] For example, the proxy server arranged on the access path to a WWW server as shown in drawing 11 is equivalent to this.

[0010] In this drawing, a proxy server 62 will be accessed to the WWW server 61 instead of a terminal 110, if the access request from a terminal 110 which is going to acquire the information on the WWW server 61 is received. Then, a proxy server 62 both returns this information to the terminal 110 which holds the information acquired by this deputy access as cache data 63 and which emitted the access request. Then, a proxy server 62 performs return processing using the cache data 63 which oneself owns, without accessing this WWW server 61, if the access request to this information is received from other terminals (graphic display abbreviation). By this, the count of access to the WWW server 61 will be mitigated.

[0011]

[Problem(s) to be Solved by the Invention] However, in the 1st

technique mentioned above, although it realizes about the cure against unloading of each server, and the backup function at the time of stopping one server, in order for a terminal utilization user to acquire required information, one of an original copy site and the mirror sites must be chosen by the user side. A user will be forced decision with which server such selection is dramatically complicated and vacant.

[0012] Moreover, in the 2nd technique mentioned above, although the count of access to the server (drawing 11 : WWW server 61) used as the information source is mitigated, the server for caches (drawing 11 : proxy server 62) cannot access this at the time of a halt of this server, but information offer to a user becomes impossible as a result.

[0013] It is in offering the server system for the Internet which can aim at backup at the time of the access distribution in a server group, and failure generating etc., without the object of this invention forcing it actuation special to a terminal utilization user in view of such a trouble.

[0014]

[Means for Solving the Problem] In the server system equipped with two or more sets of the servers which perform information dispatch through the Internet, and the management server which manages each of the server of these two or more bases when this invention for attaining the above-mentioned object was caused like 1 voice The server group which consists of said two or more sets of servers to

which a part of server [at least] of a base disseminates two or more same information is constituted. Said management server The host name given to said server group, and each address for a communication link of two or more sets of said servers which constitute said server group, The managed table on which the selection information used in case the target address for a communication link is chosen from these addresses for a communication link was stored, If the address for a communication link corresponding to said host name is required through said Internet A selection means to choose any of the address for a communication link stored in said managed table they are according to the content of said selection information, The server system characterized by having a name service means to transmit the address for a communication link which said selection means chose to the demand origin of said address for a communication link is offered.

[0015] In the server system which was equipped with two or more sets of the servers which perform information dispatch through the Internet, and the management server which manages each of the server of these two or more bases according to the mode of others of this invention for attaining the above-mentioned object The server group which consists of said two or more sets of servers to which a part of server [at least] of a base disseminates two or more same information is constituted. Said management server The identifier for server groups which expresses said server group to a meaning on said Internet, Said identifier for servers which constitutes

said server group from on said Internet and which expresses two or more each of the server of a base to a meaning, The managed table on which the selection information used in case the target identifier for servers is chosen from these identifiers for servers was stored, If the access request using said identifier for server groups is received through said Internet A selection means to choose any of said identifier for servers they are according to the content of said selection information, While transmitting the information which performed access to said server using the identifier for servers which said selection means chose, and was acquired by this access to the demand origin of said access request This information is held temporarily and the server system characterized by having a deputy access means to transmit the information concerned which holds again the access request using said identifier for server groups to the reception beam case to the demand origin of the access request concerned is offered.

[0016]

[Embodiment of the Invention] Hereafter, 1 operation gestalt of this invention is explained.

[0017] The whole server system configuration of this operation gestalt is shown in drawing 1 .

[0018] This server system has two or more information dispatch server groups (information dispatch server groups 1091, 1092, --, 109N) which perform information dispatch through the Internet 111, and the management server 101 which manages access from a terminal

performed to these information dispatch server groups. The information dispatch server groups 1091, 1092, --, 109N consist of two or more sets of servers, respectively. The information dispatch server group 1091 consists of information dispatch servers 10911-1091N which have the same information. About other information dispatch server groups as well as this, each information dispatch server owns information peculiar to the server group. In addition, what is necessary is just to add suitably the server which has the same information, when access to a specific information dispatch server group increases. The management information for managing the information dispatch server groups 1091, 1092, --, 109N is stored in the name database (name DB) 103. The DS of a name DB103 is shown in drawing 2 .

[0019] The two storing fields, the server name field 21 and the server address field 22, exist in a name DB103. With this operation gestalt, beforehand, the host name is given to information dispatch server groups [1091, 1092, -- 109N] each, and these host names are stored in the server name field 21. So to speak, this host name is a host name representing each information dispatch server of that information dispatch server group, and decides to call this host a representation host name for convenience hereafter.

[0020] On the other hand, all the information dispatch servers prepared in this server system have the server address of a proper, respectively, and are stored in the server address field 22 in the form whose server addresses of these settled for every information

dispatch server group. "-" of the server name field 21 means that it is the same as the content stored in the column on it. The address of an information dispatch server is the address for a communication link used in case it communicates on the Internet 111, and, specifically, is an IP address for performing a TCP/IP communication link.

[0021] In addition, the representation host name given to information dispatch server groups [1091, 1092, --, 109N] each does not change to the usual host name, i.e., the host name which becomes settled uniquely on DNS (Domain Name System), at all.

[0022] Here, a domain (Domain) is explained briefly.

[0023] When construction of the Internet starts, in order to identify each computer connected to the network, unitary management of the host name and IP address was carried out in the predetermined organization. The number of computers was [hundreds of] of order at the beginning, and it was also possible to have registered the host name and IP address of all computers into one file.

[0024] However, the limitation of this approach became clear immediately with network amplification. While the new registration and updating which are performed one after another take immense time amount, the rating of a managing agency increases and settlement has stopped namely, sticking as a result.

[0025] The domain was the concept proposed in order to solve such a problem, and it was introduced in order to carry out naming of many computers which exist all over the world to a meaning.

[0026] The domain space of the Internet is managed hierarchical, and DNS is designed so that this administrative information may be matched. In the Internet, a DNS server with the information on each domain is distributed, and it is managed hierarchical like domain space. The management server of this operation gestalt can also be considered to be a kind of this DNS server.

[0027] In drawing 2 , the representation host name (serv.bbb.aaa) which shows the information dispatch server group 1091, and the representation host name (other.bbb.aaa) which shows the information dispatch server group 1092 are illustrated. All are registered into the server group control table 105 also about the representation host name of information dispatch server groups other than these. The server addresses including the server address 192.128.16.1 of the information dispatch server 10911 of each information dispatch server which constitutes the information dispatch server group 1091 are stored in the form matched with the representation host name (serv.bbb.aaa) of the information dispatch server group 1091. The server address of each information dispatch server which constitutes the information dispatch server group 1092 is stored in the form matched with the representation host name (other.bbb.aaa) of the information dispatch server group 1092.

[0028] The name service section 102 in which the management server 101 performs the response processing in response to the name service demand from the outside, The name DB retrieval processing section

106 which is in the name service section 102 and performs retrieval and an extract of applicable information from a name DB103, From the name service section 102, the server group information 107 which it is as a result of [of the name DB103 by the name DB retrieval processing section 106] retrieval Reception, One server is chosen from the server group which the server group information 107 shows, and it has the server selection section 104 notified to the name service section 102 by making the result into the selection server information 108, and the server group control table 105 which the server selection section 104 refers to in case a server is chosen.

[0029] The management server 101 which has the above configuration exists on same LAN (Local Area Network)112 as the information dispatch server groups 1091-109N, and this LAN112 has connected it to the Internet 111. The countless terminal has accessed the Internet 111 and each user can access each information dispatch server group of this server system using each terminal. In addition, the terminal 110 is typically illustrated by drawing 1 .

[0030] Next, the case where the information dispatch server group 1091 is accessed from a terminal 110 is taken for an example, and actuation of this server system is explained. Here, the talk is advanced as that to which the user knows beforehand the representation host name given to the information dispatch server group 1091.

[0031] A user inputs the representation host name (serv.bbb.aaa) of the information dispatch server group 1091 into a terminal 110

first. A terminal 110 asks the management server 101 the server address corresponding to a representation host name in response to this alter operation. Specifically, the information which shows a representation host name (serv.bbb.aaa) is sent to the management server 101 in this case. So far, it is the same as the case where a name service is required of the conventional DNS server.

[0032] The name service section 102 of the management server 101 will start the name DB retrieval processing section 106, if the information which shows, the address inquiry (serv.bbb.aaa), i.e., the representation host name, from a terminal 110, is received. The name DB retrieval processing section 106 performs name DB retrieval processing shown in drawing 5.

[0033] As shown in this drawing, the name DB retrieval processing section 106 searches a name DB103 for the representation host name sent from the terminal 110 to a key first (step 51 (S51)). Thereby, all the server addresses corresponding to the representation host name used as the key are extracted. Here, the server address (192.128.16.1, 192.128.16.4, --) corresponding to a representation host name (serv.bbb.aaa) is extracted. Next, the name DB retrieval processing section 106 notifies all the found server addresses to the server selection section 104 as server group information 107 with a previous representation host name (S52). The server selection section 104 chooses the one optimal server for making information offer perform with reference to the server group control table 105 according to the content. Suppose that the detail

of this selection processing is mentioned later. The selected host name and selected server address of a server are sent to the name service section 102 as selection server information 108. The name service section 102 processes this server address as a retrieval result, and transmits this server address to a terminal 110 as response processing of an address inquiry demand after that (S53).

[0034] The terminal 110 to which the server address was notified can perform information access now from the management server 101 directly to the server after this.

[0035] Thus, since the optimal server for performing information dispatch out of the information dispatch server group whenever a certain server group from information is accessed is chosen according to this operation gestalt, access from a terminal will be distributed within the information dispatch server group as a result.

[0036] In addition, at the terminal 110, although not explained especially, it succeeds in configuration beforehand so that the representation host name sent out from this terminal 110 may reach the management server 101 through DNS which is not illustrated direct or here. About this, since it is usually carried out on DNS, it omits for details.

[0037] Moreover, with the above-mentioned operation gestalt, although the information dispatch server group and the management server are connected to the same LAN, this does not mean that an information dispatch server group must exist on the same LAN as a

management server. For example, each server of an information dispatch server group may be connected to the location of the arbitration on the Internet.

[0038] Below, the operation gestalt of others of this invention is explained.

[0039] The whole server system configuration of this operation gestalt is shown in drawing 3 .

[0040] The management server 301 of this server system plays the role of the so-called proxy server which performs information access of as opposed to a WWW server for access to a WWW server from a terminal, and the cash advance of the access information instead of a direct receptionist and a terminal in access to the WWW (World Wide Web) server which disseminates various multimedia information by the Internet.

[0041] 1091-109N are information dispatch server groups which have the same information and which consist of server two or more bases among drawing. Here, the information dispatch server group 1091 consists of two or more sets of WWW servers, and the information dispatch server group 1092 consists of two or more sets of FTP servers.

[0042] In the management server 301, 302 responds to information access from a terminal. The deputy access section which performs deputy access to an information dispatch server, return to the terminal of access information, and momentary maintenance (cache processing) of the access information, and 306 The server list

retrieval processing section which is in the deputy access section 302 and performs retrieval and an extract of applicable information from the server list 303, and 304 From the deputy access section 302, the server group information 307 which it is as a result of [of the server list 303 by the server list retrieval processing section 306] retrieval Reception, It is the server group control table which chooses one server from the server groups which the information shows, and the server selection section which notifies the result to the deputy access section 302 as selection server information 308, and 305 have in the server selection section 304, and is referred to in the case of selection processing of a server.

[0043] The DS of the server list 303 is shown in drawing 4 .

[0044] As shown in drawing 4 , the two storing fields, the representation server name field 41 and the server name field 42 in a server group, exist in the server list 303. With this operation gestalt as well as the above-mentioned, the host name (representation host name) is given to information dispatch server groups [1091, 1092, --, 109N] each, and these representation host names are stored in the server name field 41. "-" of the representation server name field 41 means that it is the same as the content stored in the column on it. Moreover, the host name of each server which constitutes the information dispatch server group shown by the representation host name stored in the representation server name field 41 is stored in the server name field 42 in a server group.

[0045] The representation host name (www.bbb.aaa) which shows the information dispatch server group 1091, and the representation host name (ftp.bbb.aaa) which shows the information dispatch server group 1092 are illustrated by drawing 4 . The host names including the host name (serv1.www-c.bbb.aaa) of the information dispatch server 10911 of each information dispatch server which constitutes the information dispatch server group 1091 are stored in the form matched with the representation host name (www.bbb.aaa) of the information dispatch server group 1091. The host name of each information dispatch server which constitutes the information dispatch server group 1092 is stored in the form matched with the representation host name (ftp.bbb.aaa) of the information dispatch server group 1092.

[0046] The information dispatch server groups 1091-109N and the access control server 301 are connected to same LAN (Local Area Network)112, respectively. This LAN112 is connected to the Internet 111. A user's terminal 110 used in case the information which the information dispatch server groups 1091-109N are disseminating comes to hand has accessed the Internet 111.

[0047] The case where the information which each server of the information dispatch server group 1091 is disseminating next comes to hand at a terminal 110 is taken for an example, and actuation of this server system is explained. Here, the talk is advanced as that to which the user knows beforehand the representation host name given to the information dispatch server group 1091.

[0048] A user inputs the representation host name (www.bbb.aaa) of the information dispatch server group 1091 into a terminal 110 first. A terminal 110 requires access to delivery and the information dispatch server group 1091 of the management server 301 for a representation host name in response to this alter operation. So far, it is the same as the case where access is required of the conventional proxy server.

[0049] The deputy access section 302 of the management server 301 will start the server list retrieval processing section 306, if a representation host name (www.bbb.aaa) is received from a terminal 110. The server list retrieval processing section 306 performs server list retrieval processing shown in drawing 6 .

[0050] The server list retrieval processing section 306 searches the server list 303 for the representation host name sent from the terminal 110 to a key first (step 61 (S61)). Thereby, all the host names corresponding to the representation host name used as the key are extracted. Here, the host name (serv1.www-c.bbb.aaa, serv4.www-c.bbb.aaa, --) corresponding to a representation host name (serv.bbb.aaa) is extracted. Next, the server list retrieval processing section 306 notifies all the found host names to the server selection section 304 as server group information 307 with a previous representation host name (S62). The server selection section 304 chooses the one optimal server for making information offer perform with reference to the server group control table 305 according to the content. Suppose that the detail of this selection

processing is mentioned later. The host name of the selected server is sent to the deputy access section 302 as selection server information 308 with said representation host name. The deputy access section 302 accesses that server using this host name. The deputy access section 302 stores in the memory in this management server 301 the information acquired by this deputy access as cache data, and returns this cache data further to a terminal 110. When the access request by the representation host name (www.bbb.aaa) is received next time, this cache data is transmitted to the access request origin concerned.

[0051] By the above actuation, a terminal 110 accesses an information dispatch server indirectly through the management server 301.

[0052] Thus, since deputy access which the management server 301 performs is performed to the optimal server at the event according to this operation gestalt, access in the server group is distributed as a result.

[0053] In the above, although two operation gestalten of the server system of this invention were explained focusing on the whole configuration and actuation, the detail of the server selection section with which a management server is equipped is explained.

[0054] Here, taking the case of the server system shown in drawing 1 , the talk is advanced for convenience.

[0055] First, actuation of the server selection section 104 at the time of using the server group control table 105 shown in drawing

7 is explained.

[0056] Although two or more server addresses which they are as a result of [which was searched considering the representation host name used as the key by name DB retrieval processing (refer to drawing 5) and this representation host name as a key] retrieval are contained in the server group information 107 sent to the server selection section 104 from a name service 102, said representation host name is stored in the server name field 1051 of drawing 7 . Moreover, two or more above-mentioned server addresses which they are as a result of retrieval are stored in the server address field 1052.

[0057] Moreover, although the host name of the server which the server selection section 104 chose is sent to a name service 102 as selection server information 108, about the server address notified for the selection server information 108, "1" (1: advice) is set as the access flag field 1053. In addition, "0" (0: un-notifying) is set up about other server addresses.

[0058] And if the server selection section 104 judges whether the representation host name which the server group information 107 shows exists in the server name field 1051 and this representation host name cannot be found in the server name field 1051 when the server group information 107 is received, it will register into the server group control table 105 all the server addresses (all server addresses contained in the server group information 107) matched with this with this representation host name. At this time, "1" is

set as that access flag field 1053 about one server address. "0" is set to other access flag fields 1053. In this drawing, "1" was set as the access flag field 1053 of a server address 192.128.16.1, and "0" is set up except this.

[0059] On the other hand, the next activity is done when the same representation host name is already stored in the server name field 1051.

[0060] For example, when the content of the server group control table 105 is already set as the condition of drawing 7 , while setting "1" as the access flag field 1053 of a server address 192.128.16.4, "0" is set as the access flag field 1053 of a server address 192.128.16.1. In addition, when "1" is set as the access flag field 1053 of a server address 192.128.16.12 located in the tail end, while setting "1" as the access flag field 1053 of a server address 192.128.16.1 located in a head line, "0" is set as the access flag field 1053 of a server address 192.128.16.12.

[0061] When the above processing is performed, "1" which is the set point is visible as if it moved in order in the access flag field 1053.

[0062] Next, the server selection section 104 is notified to the name service section 102 by making into the selection server information 108 the server address stored in the server address field 1052 of the line which set up "1" in the access flag field 1053.

[0063] Thus, if the server to which information dispatch is made

to perform is chosen one after another, access from a terminal will come to be distributed almost uniformly to each server in the information dispatch server group.

[0064] Actuation of the server selection section 104 at the time of next using the server group control table 105 shown in drawing 8 is explained.

[0065] Here, the server to access is chosen according to the load of each server which constitutes an information dispatch server group.

[0066] In drawing 8 , about the same component, the same number is attached and explanation is omitted.

[0067] This server group control table 105 is replaced with the previous access flag field 1053, and the description is in a point equipped with the load level field 1054. The information which shows the load profile initiation of each server is stored in the load level field 1054.

[0068] moreover, new in the selection section 104 -- the what server load Monitoring Department 113 is formed. The server load Monitoring Department 113 supervises periodically the load profile initiation of the server (namely, server by which the server address is stored in the address field 1052) registered into the server group control table 105, and sets the load level as the load level field 1054. The load level is beforehand divided into several steps (here three-stage of load levels A, B, and C). The server load Monitoring Department 113 transmits a monitor command periodically

to the server of the server address stored in the server address field 1052, totals the response time, divides the result for any of load level A/B/C being, and, specifically, sets it as the load level field 1054. Here, about the server which does not answer at all to the above-mentioned monitor command, nothing is set up to the load level field 1053. When no servers answer, it is expected that the server is in the idle state by generating of a failure among a maintenance service.

[0069] And the server selection section 104 chooses what has the lightest load from the servers registered into the server group control table 105 in the server group information 107 at the carrier beam event, and notifies it to the name service section 102 by making the address of this server etc. into the selection server information 108. Under the present circumstances, suppose that it excepts from a selection branch beforehand about the server which is not set as the load level field 1053 at all.

[0070] If such processing is performed, access from a terminal will come to be performed to a server with the lightest load at the event out of the target information dispatch server group. Moreover, access to the server which is not working will also be forbidden.

[0071] In addition, when choosing the server accessed according to the load of each server, server SHITEMU may be constituted as shown in drawing 9 .

[0072] In this drawing, it acts to each information dispatch server as the monitor of the load of the server, and the load monitor 115

which reports the load profile initiation to the management server 101 is formed in it. The management server 101 collects the load situation reports from each server, and is equipped with the load profile initiation collection section 114 stored in the server group control table 105 by making the result into a load level. Here, about a server without a report, a load level is not set up from the load monitor section 115.

[0073] And if the server group information 107 is received, the server selection section 104 will discover that to which load level information is set out of the server registered into the server group control table 105, will choose from that inside what has the lightest load, and will notify it to the name service section 102 by making the server address of this server into the selection server information 108.

[0074] Thus, if it acts to the server itself as the monitor of the load profile initiation of each server, more exact load information can be acquired.

[0075] As mentioned above, although the various operation gestalten of the server selection section with which a management server is equipped were explained, these are applicable also to the server system shown in drawing 3 .

[0076]

[Effect of the Invention] According to this invention, as cures against backup at the time of failure generating etc., since the optimal server for making information dispatch perform is

automatically chosen even if it is the case where two or more sets of the servers which have the same information are prepared, a user can get required information, without being conscious of existence of two or more sets of servers, in order to distribute access from a user.

[0077] Moreover, according to this invention, since the server to which the optimal server for making information dispatch perform is not working, corresponding to the load profile initiation of each server is chosen as outside of an object, a user can access it in the best environment.

[0078] Even if access load moreover increases, the engine-performance reinforcement according to access load is attained by offering this invention only by adding the server which has the same information.

(End of Document)

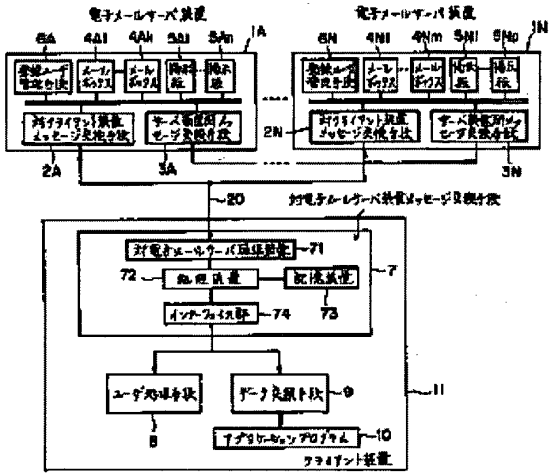
ELECTRONIC MAIL MULTIPLEXING SYSTEM AND COMMUNICATION CONTROL METHOD IN THE SYSTEM

Publication number: JP7066829
Publication date: 1995-03-10
Inventor: SAITO HIROSHI
Applicant: KAWASAKI STEEL CO
Classification:
- international: ***G06F13/00; H04L12/54; H04L12/58;*** **G06F13/00; H04L12/54; H04L12/58;** (IPC1-7): H04L12/54; G06F13/00; H04L12/58
- European:
Application number: JP19930162529 19930630
Priority number(s): JP19930162529 19930630

Report a data error here

Abstract of **JP7066829**

PURPOSE:To distribute loads between multiplexed electronic mail server equipments by continuing service when at least one of plural mail server equipments is in operation. CONSTITUTION:Each electronic mail server equipment is provided with a function copying a content of a mail box or an electronic bulletin board provided to each of multiplexed electronic mail server equipments 1A-1N. A message exchange means 7 of a counter electronic mail server equipment in a client equipment 11 automatically selects one of the electronic mail server equipments in operation and connects the both. Furthermore, the client equipment selects the electronic mail server equipment with least load based on the information of the electronic mail server equipment.



Data supplied from the *esp@cenet* database - Worldwide

(19)日本国特許庁（ＪＰ）　(12) **公開特許公報（A）**　(11)特許出願公開番号

特開平7－66829

(43)公開日　平成７年(1995)３月10日

(51)Int.Cl.[6]	識別記号	庁内整理番号	ＦＩ	技術表示箇所
Ｈ０４Ｌ　12/54				
Ｇ０６Ｆ　13/00	３５１　Ｇ	7368－5Ｂ		
Ｈ０４Ｌ　12/58				
		8732－5Ｋ	Ｈ０４Ｌ　11/ 20　１０１　Ｂ	

審査請求　未請求　請求項の数５　ＯＬ　（全　８　頁）

(21)出願番号　特願平5－162529

(22)出願日　平成５年(1993)６月30日

(71)出願人　000001258
川崎製鉄株式会社
兵庫県神戸市中央区北本町通１丁目１番28号

(72)発明者　斎藤　洋
東京都千代田区内幸町２丁目２番３号　川崎製鉄株式会社東京本社内

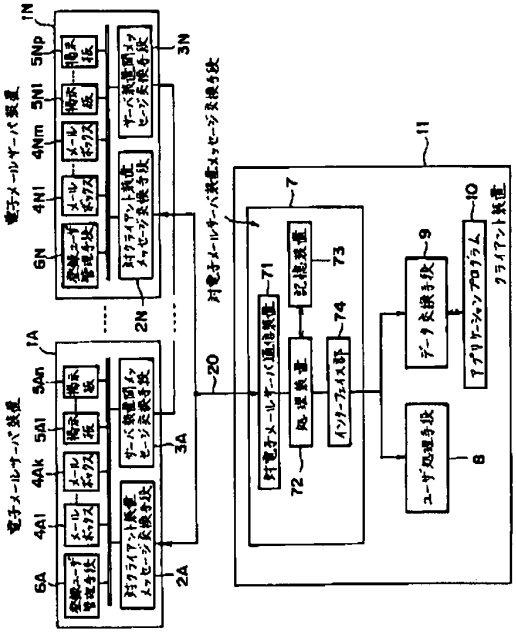
(74)代理人　弁理士　谷　義一　（外１名）

(54)【発明の名称】　電子メール多重化システムおよび該システムにおける通信制御方法

(57)【要約】

【構成】　多重化された電子メールサーバ装置（１Ａ－１Ｎ）がそれぞれ有する、メールボックス（または、電子掲示板）の内容を、相互に複写する機能を、各電子メールサーバ装置にもたせ、クライアント装置（１１）内の対電子メールサーバ装置メッセージ交換手段（７）が、稼働中の電子メールサーバ装置の１台を自動的に選択して、両者を接続するようにした。また、電子メールサーバ装置からの情報に基づいて、最も負荷の少ない電子メールサーバ装置をクライアント装置が選択する。

【効果】　複数の電子メールサーバ装置のうち、少なくとも１台が稼働状態であれば、サービスが継続される。多重化された電子メールサーバ装置間の負荷分散が可能である。



1

【特許請求の範囲】

【請求項1】　多重化された電子メールサーバ装置とクライアント装置とがネットワークを介して接続された電子メールシステムにおいて、

前記電子メールサーバ装置は、

メールボックスまたは電子掲示板等のメッセージ蓄積手段と、

登録されたユーザの管理をする登録ユーザ管理手段と、

前記クライアント装置とのメッセージの送受信を行う対クライアント装置メッセージ交換手段と、

前記電子メールサーバ装置間でのメッセージの交換を行うサーバ装置間メッセージ交換手段とを備え、

前記クライアント装置は、

使用者またはアプリケーションプログラムとのインターフェイスをとる手段と、

前記電子メールサーバ装置とのメッセージの送受信を行う対電子メールサーバ装置メッセージ交換手段とを備え、

前記サーバ装置間メッセージ交換手段は、複数の電子メールサーバ装置間で、前記メッセージ蓄積手段の内容を複写し、

前記対電子メールサーバ装置メッセージ交換手段は、前記電子メールサーバ装置からの情報に基づいて、前記電子メールサーバ装置の中から、稼働中の装置を通信相手として選択することによって、電子メールシステムの無停止を実現することを特徴とする電子メール多重化システム。

【請求項2】　請求項1において、前記対電子メールサーバ装置メッセージ交換手段は、前記電子メールサーバ装置からの情報に基づいて、負荷の少ない電子メールサーバ装置を選択することを特徴とする電子メール多重化システム。

【請求項3】　多重化された電子メールサーバ装置とクライアント装置とがネットワークを介して接続された電子メールシステムにおいて、

前記電子メールサーバ装置間でメッセージの交換を行い、複数の電子メールサーバ装置間で、前記メッセージを複写する第1の過程と、

前記クライアント装置が、前記電子メールサーバ装置からの情報に基づいて、前記電子メールサーバ装置の中から、稼働中の装置を通信相手として選択する第2の過程とを具備することを特徴とする電子メール多重化システムにおける通信制御方法。

【請求項4】　請求項3において、前記第2の過程は、

前記クライアント装置が前記電子メールサーバ装置に対して自局のＩＤを一斉送信する過程と、

該送信に応じた複数の電子メールサーバ装置の中から一つの電子メールサーバ装置を選択する過程と、

選択した電子メールサーバ装置に接続要求する過程とを具備することを特徴とする電子メール多重化システムにおける通信制御方法。

2

【請求項5】　請求項4において、前記クライアント装置が、通信相手の電子メールサーバ装置の故障を検出した場合には、前記一斉送信に応じた複数の電子メールサーバ装置の中から稼働中の他の装置を選択して、接続要求を送信することを特徴とする電子メール多重化システムにおける通信制御方法。

【発明の詳細な説明】

【０００１】

【産業上の利用分野】本発明は、多重化された電子メールサーバ装置を備えた電子メール多重化システムおよび該システムにおける通信制御方法に関する。

【０００２】

【従来の技術】従来の電子メールシステムとしては、特開平４－２７３６３６号公報に記載された分散掲示板方式が知られている。これは、電子メールシステムにおいて、複数のサーバ装置に掲示板のコピーを配置することにより、掲示物の読み出し要求が１つのサーバ装置に集中することを防止し、応答性を向上させようとしたものである。

【０００３】この従来の電子メールシステムにおいては、電子メールサーバ装置とクライアント装置とが、ネットワークを介して、固定的に接続されていた。すなわち、特定の利用者が、電子メールシステムから必要なデータを得る場合は、特定の１台の電子メールサーバ装置からそのデータを得ていた。なお、この種のデータを利用するものとしては、一般に、利用者とアプリケーションプログラムとがあるが、以下の説明においては、これらを総称して、「サービス受益対象」と呼ぶこととする。

【０００４】

【発明が解決しようとする課題】従来の電子メールシステムにおいては、特定のサービス受益対象と電子メールサーバ装置とが固定的に接続されているため、自分が接続された電子メールサーバ装置が停止した場合、サービス受益対象は、この電子メールシステムを利用することができなかった。このような事態を避けるためには、電子メールシステムの無停止サービスを実現しなければならない。具体的には、電子メールサーバ装置にフォールトトレラント型のコンピュータを適用したり、複数の電子メールサーバ装置を高度に結合し、多重化する方法がとられていた。これらの装置を実現するには、高度な制御技術と、多大な費用が要求された。

【０００５】さらに、同一のネットワーク上に複数の電子メールサーバ装置が稼働していても、特定の電子メールサーバ装置に、アクセスが集中してしまう可能性があった。

【０００６】本発明は、無停止で稼働する安価な電子メール多重化システムを提供することを目的とする。また、多重化された電子メールサーバ装置間の負荷分散を

3

可能にしたシステムを提供することを目的とする。

【0007】

【課題を解決するための手段】上記課題を解決するために、本発明は、多重化された電子メールサーバ装置に、それぞれが有するメールボックス（または、電子掲示板）の内容を、相互に複写する機能をもたせるとともに、クライアント装置が、稼働中の電子メールサーバ装置の1台を自動的に選択して接続するようにした。すなわち、本発明では、電子メールサーバ装置の切り替えを、クライアント装置が行うようにした。

【0008】また、電子メールサーバ装置からの情報に基づいて、最も負荷の少ない電子メールサーバ装置をクライアント装置が選択することによって、複数の電子メールサーバ装置間の負荷のバランスをとっている。

【0009】

【作用】本発明によれば、複数の電子メールサーバ装置のうち、少なくとも1台が稼働状態であれば、クライアント装置内に配置されている対電子メールサーバ装置メッセージ交換手段が、稼働中の電子メールサーバ装置の中の1台を自動的に選択して接続を行う。

【0010】さらに、クライアント装置が接続する際に、電子メールサーバ装置側からの情報に基づいて、負荷の最も少ない電子メールサーバ装置に接続することによって、多重化された電子メールサーバ装置間の負荷分散を可能にした。

【0011】

【実施例】次に、本発明の実施例を詳細に説明する。

【0012】図1は、本発明の一実施例を示す図である。

【0013】図1において、電子メールサーバ装置1A－1Nは、メッセージを蓄積して、要求に応じてメッセージを要求元に転送する装置である。また、クライアント装置11は、サービス受益対象が、電子メールサーバ装置1A－1Nにサービスを要求するための装置である。

【0014】電子メールサーバ装置1A－1Nは、それぞれ、対クライアント装置メッセージ交換手段2A－2N、サーバ装置間メッセージ交換手段3A－3N、メールボックス4A1－4Ak－4N1－4Nm、掲示板5A1－5An－5N1－5Np、および登録ユーザ管理手段6A－6Nを備えている。

【0015】ここで、各メールボックスおよび掲示板は、メッセージ蓄積手段として機能するもので、サーバ装置間メッセージ交換手段3A－3Nを介して、相互に接続されている。登録ユーザ管理手段6A－6Nは、電子メールサーバ装置1A－1Nに接続を要求してきたサービス受益対象の認証を行うもので、内部にユーザIDとパスワードとを格納している。

【0016】クライアント装置11は、対電子メールサーバ装置メッセージ交換手段7と、ユーザ処理手段8

4

と、データ交換手段9と、アプリケーションプログラム10とを備えている。ここで、対電子メールサーバ装置メッセージ交換手段7は、各電子メールサーバ装置との通信を制御する対電子メールサーバ通信装置71と、処理装置72と、プログラムおよび各種データを格納する記憶装置73と、インターフェイス部74とを有している。対電子メールサーバ通信装置71は、処理装置72からの要求によって、図2および図3に示す通信制御を行う。また、インターフェイス部74は、処理装置72とユーザ処理装置8およびデータ交換手段9とのインターフェイスをとる。

【0017】ユーザ処理手段8は、利用者からのメッセージや指示を受け取り、電子メールサーバ装置1A－1Nに対して、メッセージを転送したり、メッセージの転送を要求したり、また転送されてきたメッセージを利用者に供給する。電子メールサーバ装置との、これらのやり取りは、クライアント装置内の対電子メールサーバ装置メッセージ交換手段7、ネットワーク20、電子メールサーバ装置内の対クライアント装置メッセージ交換手段2A－2Nを介して行われる。

【0018】アプリケーションプログラム10は、クライアント装置11内に、必要に応じて配置されたもので、データ交換手段9および対電子メールサーバ装置メッセージ交換手段7を通して、電子メールサーバ装置1A－1Nと接続され、メッセージ交換を行う。

【0019】次に、電子メールサーバ装置1（1A－1N）の多重化について説明する。

【0020】ある電子メールサーバ装置内のメールボックス4および掲示板5は、サーバ装置間メッセージ交換手段3を通して、他の電子メールサーバ装置内のメールボックス4および掲示板5にコピーされる。このコピーは、ネットワークの構成や、メッセージの転送頻度などに応じて決められる。具体的には、夜間や休日など、装置が比較的すいている間に、一括コピーする方法や、メッセージがいずれかの電子メールサーバ装置に到着した時点で、ただちに、他の電子メールサーバ装置のメールボックスや掲示板にコピーする方法などがある。なお、コピーの数は、必ずしも同一である必要はなく、メッセージの重要性に応じて、変えてもよい。

【0021】次に、図2および図3を参照して、本実施例の動作を説明する。

【0022】図2は、接続時の手順を示すものであり、サービス受益対象（ユーザまたはアプリケーションプログラム）、対電子メールサーバ装置メッセージ交換手段7、および電子メールサーバ装置の3つに分けて、動作を示している。

【0023】サービス受益対象は、電子メールサーバ装置1からのサービスを受けるために、自身のIDおよびパスワードを、クライアント装置11に指示する（ステップS1－S4）。クライアント装置11のユーザ処理

5

手段8またはデータ交換手段9は、対電子メールサーバ装置メッセージ交換手段7に対して、ＩＤとパスワードとを引き渡した上で、電子メールサーバ装置1への接続を依頼する。このとき、対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象のＩＤとパスワードとを内部的に記憶しておく。

【００２４】次に、対電子メールサーバ装置メッセージ交換手段7は、接続可能なすべての電子メールサーバ装置1に対して、サービス受益対象のＩＤと特定の指示とを、一斉送信する（ステップＳ５）。ここで、特定の指示とは、上記ＩＤをもつサービス受益対象にサービス提供可能な電子メールサーバ装置は、その電子メールサーバ装置の識別コードを返送するように、との指示である。

【００２５】一斉送信を受信した電子メールサーバ装置1Ａ－1Ｎは、受信したサービス受益対象のＩＤと、自分が管理する登録ユーザ管理手段6Ａ－6Ｎ内のデータとを照合し、一致するものがあれば、自分の識別コードを、対クライアント装置メッセージ交換手段2Ａ－2Ｎを経由して、問い合わせ元の対電子メールサーバ装置メッセージ交換手段7に返す（ステップＳ６－Ｓ９）。こうして、このサービス受益対象にサービスを提供できる電子メールサーバ装置の識別コードが、対電子メールサーバ装置メッセージ交換手段7に受信される（ステップＳ１０）。

【００２６】対電子メールサーバ装置メッセージ交換手段7は、サービスを提供できる電子メールサーバ装置のリストを作成し（ステップＳ１１）、接続すべき電子メールサーバ装置を選択する（ステップＳ１３）。この場合、各電子メールサーバ装置に対して、現在の負荷状況などの問い合わせをして（ステップＳ１２）、電子メールサーバ装置を決定してもよい。一方、サービス提供可能な電子メールサーバ装置が存在しない場合、対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象に対して、サービス提供ができない旨の通知を行う（ステップＳ１４－Ｓ１５）。サービス受益対象は、この通知を受信して、接続が失敗であったことを認識し、処理を終了する（ステップＳ１６）。

【００２７】利用可能な電子メールサーバ装置が複数存在する場合、対電子メールサーバ装置メッセージ交換手段7は、あらかじめ定められた手順で、一台の電子メールサーバ装置を選択する。具体的には、次のような方法が考えられる。

【００２８】（1）　サービス可能なすべての電子メールサーバ装置に対して、処理負荷に関する問い合わせを行い（ステップＳ１２）、その時点で負荷が最も少ない電子メールサーバ装置を選択する。

【００２９】（2）　対電子メールサーバ装置メッセージ交換手段7に、最初に識別コードを返した電子メールサーバ装置を選択する。

6

【００３０】電子メールサーバ装置の選択が完了すると、対電子メールサーバ装置メッセージ交換手段7は、選択した電子メールサーバ装置に対して、ログイン要求、ユーザＩＤ、およびパスワードを送信する（ステップＳ１７）。電子メールサーバ装置は、これらを受信し（ステップＳ１８）、接続が可能か否かを確認のうえ（ステップＳ１９）、接続を実行する（ステップＳ２０）。接続完了後、あるいは接続が不可能な場合、電子メールサーバ装置は、対電子メールサーバ装置メッセージ交換手段7に対して、結果を通知する（ステップＳ２１）。

【００３１】対電子メールサーバ装置メッセージ交換手段7は、接続結果を受信し（ステップＳ２２）、接続が成功したときには、サービス開始をサービス受益対象に通知する（ステップＳ２３－Ｓ２４）。サービス受益対象は、このサービス開始通知を受信し、次の段階（サービス要求）に入る。一方、接続が失敗に終った場合は、ステップＳ１３に戻って、他の電子メールサーバ装置を選択し、この電子メールサーバ装置に接続要求を行う。指定した電子メールサーバ装置との接続に再度失敗した場合は、この電子メールサーバ装置を、サービス可能な電子メールサーバ装置のリストから外して、他の電子メールサーバ装置にもう一度接続要求を行う。

【００３２】このようにして、電子メールサーバ装置との接続に成功するか、あるいは、サービス可能な電子メールサーバ装置がなくなるまで、ステップＳ１３－Ｓ２３の手順を繰り返す。

【００３３】図3は、正常時のサービス要求の動作と、電子メールサーバ装置に障害が発生した場合の動作とを示すフローチャートである。

【００３４】電子メールサーバ装置との接続に成功すると、サービス受益対象は、対電子メールサーバ装置メッセージ交換手段7に対して、サービス要求を送信する（ステップＳ３１）。このサービス要求を受信すると、対電子メールサーバ装置メッセージ交換手段7は、電子メールサーバ装置に対して、サービス要求を送信する（ステップＳ３２）。サービス要求を受信した電子メールサーバ装置は、サービスを実行し、サービス内容を対電子メールサーバ装置メッセージ交換手段7に返送する（ステップＳ３３－Ｓ３５）。

【００３５】サービス内容を受信した対電子メールサーバ装置メッセージ交換手段7は、サービス内容をサービス受益対象に転送する（ステップＳ３６－Ｓ３７）。サービス受益対象はこのサービス内容を受領して処理を終了する（ステップＳ３８）。電子メールサーバ装置との接続に成功し、電子メールサーバ装置からのサービスの提供が可能になった場合、電子メールサーバ装置と対電子メールサーバ装置メッセージ交換手段7は、サービス受益対象が電子メールサービスが不要となるまで、その接続を継続する。もし、何らかの原因で、接続中の電子

メールサーバ装置がサービスを提供できなくなった場合は、次の障害発生時のフローにしたがって、サービスの再要求が行われる。

【００３６】対電子メールサーバ装置メッセージ交換手段７が、電子メールサーバ装置に対してサービス要求を送信した後、あらかじめ定めた時間、サービス内容が返送されない場合、対電子メールサーバ装置メッセージ交換手段７は、サービス内容の受信に失敗したものと判断し（ステップＳ４１）、電子メールサーバ装置に対してサービスの再要求を行う。すなわち、ステップＳ１３－Ｓ２３において、電子メールサーバ装置を選択し、選択した電子メールサーバ装置に対して、サービス再要求を送信する（ステップＳ３２）。なお、この手順は、図２の接続時の手順と同様である。この手順は、サービス受益対象に何ら通知せずに行われる。すなわち、サービスを提供する電子メールサーバ装置が障害を起こすたびに、対電子メールサーバ装置メッセージ交換手段７は、他のサービス可能な電子メールサーバ装置とクライアント装置とを順次接続し、サービスの要求およびサービス内容の転送を実行する（ステップＳ３２－Ｓ３８）。こうして、無停止の電子メールシステムを実現することができる。

【００３７】ただし、サービス可能な電子メールサーバ装置が無くなった場合には、対電子メールサーバ装置メッセージ交換手段７は、サービス受益対象に対して、サービス提供不能通知を送信し（ステップＳ１５）、サービス受益対象は、電子メールサーバ装置が障害である旨の報告を行う（ステップＳ１６ａ）。

【００３８】【発明の効果】以上説明したように、この発明は、電子メールサーバ装置の切り替えをクライアント装置側で行うことによって、安価な電子メールシステムを実現することができる。また、処理負荷の最も軽い電子メールサーバ装置を選択する方式を使用することによって、負荷分散を実現することができる。

【図面の簡単な説明】

【図１】この発明による電子メール多重化システムの実施例の構成を示すブロック図である。

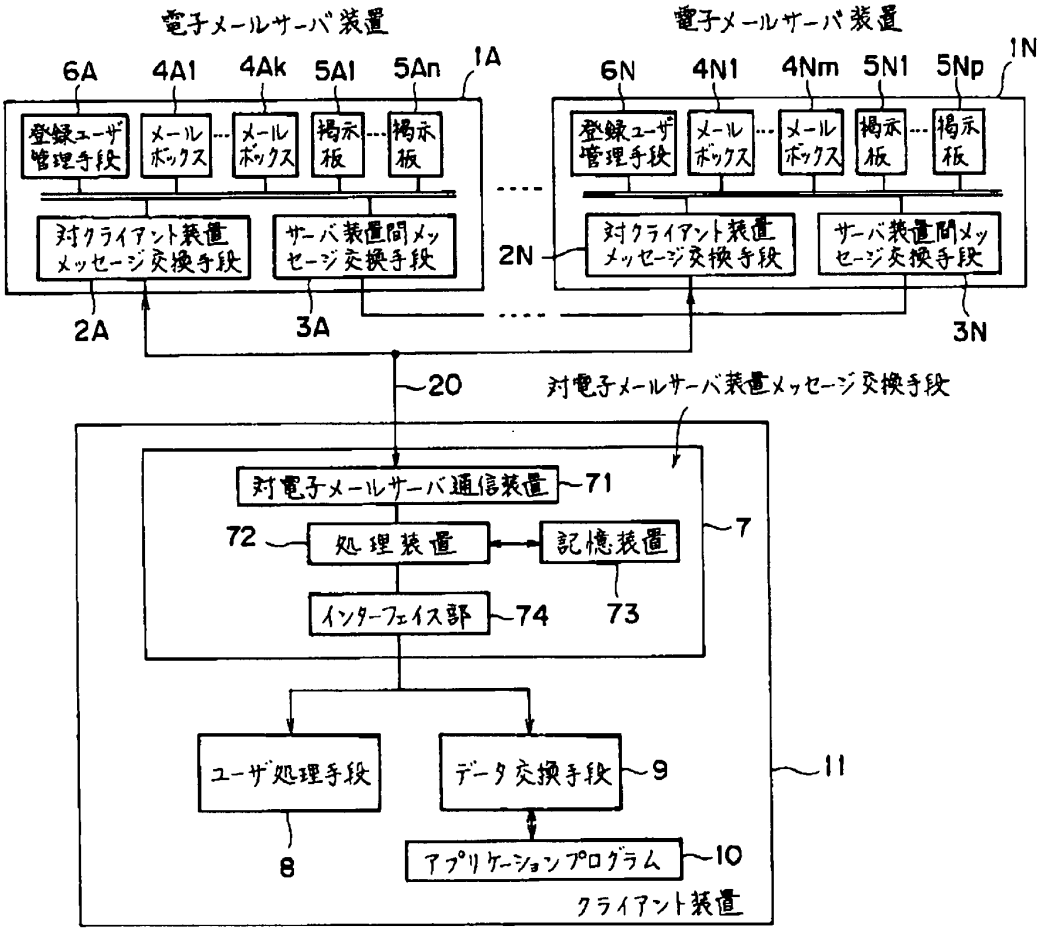
【図２】上記実施例の接続時の動作を示すフローチャートである。

【図３】上記実施例のサービス要求時の動作を、正常時と、障害発生時とに分けて示すフローチャートである。

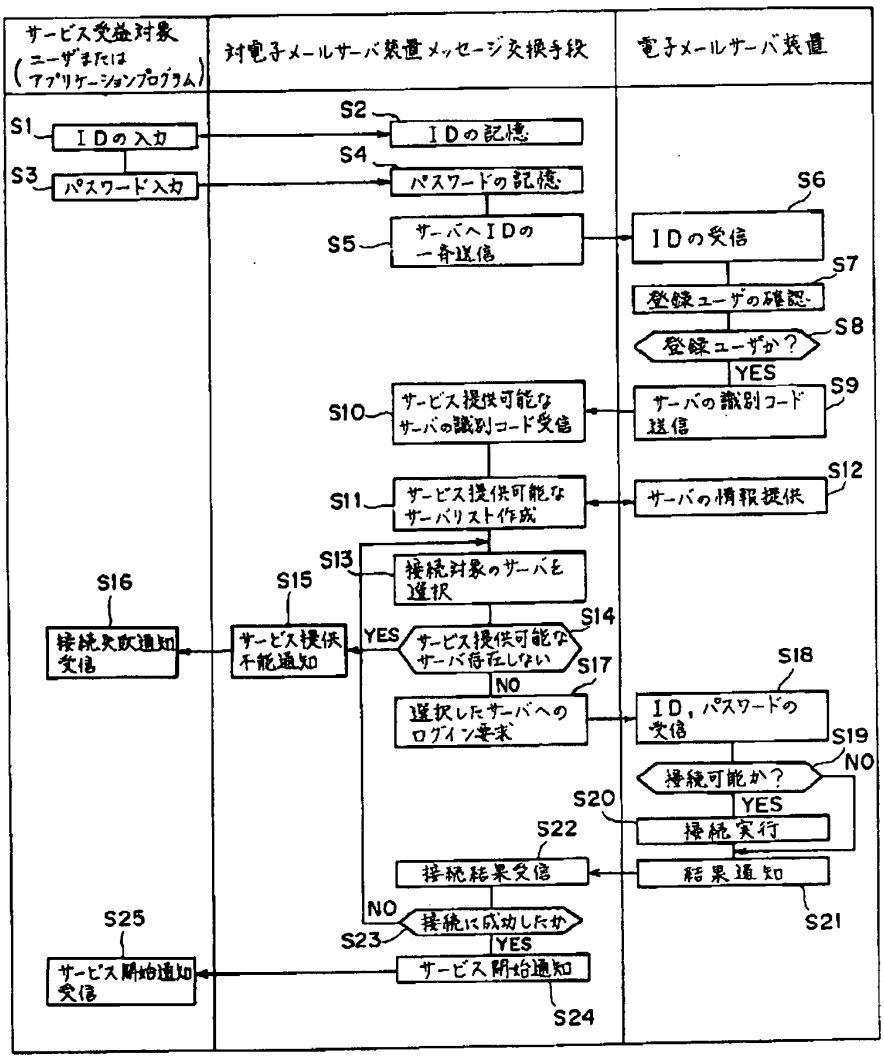
【符号の説明】

１Ａ，１Ｎ　電子メールサーバ装置
２Ａ，２Ｎ　対クライアント装置メッセージ交換手段
３Ａ，３Ｎ　サーバ装置間メッセージ交換手段
４Ａ１－４Ａｋ，４Ｎ１－４Ｎｍ　メールボックス
５Ａ１－５Ａｎ，５Ｎ１－５Ｎｐ　掲示板
６Ａ－６Ｎ　登録ユーザ管理手段
７　対電子メールサーバ装置メッセージ交換手段
８　ユーザ処理手段
９　データ交換手段
１０　アプリケーションプログラム
２０　ネットワーク
７１　対電子メールサーバ通信装置
７２　処理装置
７３　記憶装置
７４　インターフェイス部

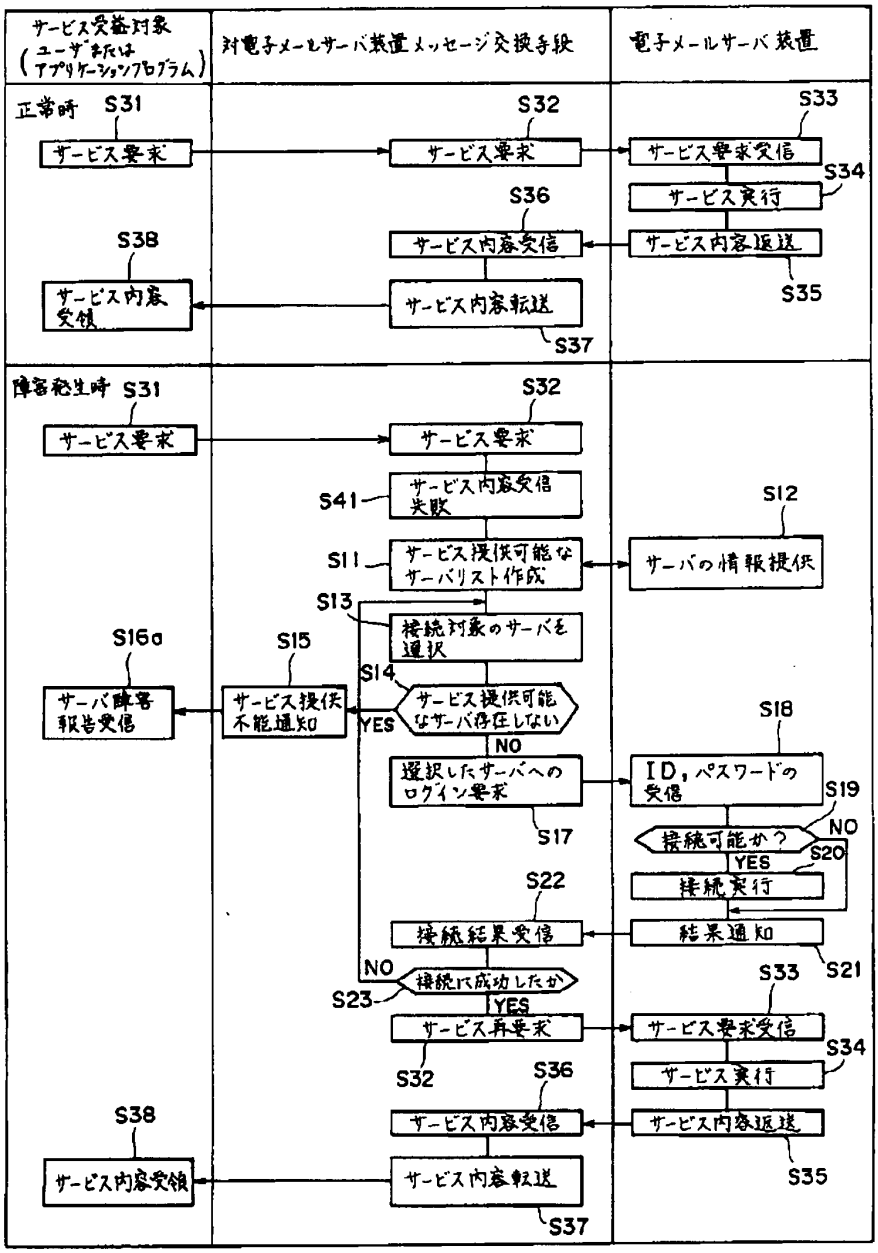
【図1】


電子メールサーバ装置
電子メールサーバ装置
6A
4A1
4Ak
5A1
5An
1A
6N
4N1
4Nm
5N1
5Np
1N
登録ユーザ管理手段
メールボックス
メールボックス
掲示板
掲示板
登録ユーザ管理手段
メールボックス
メールボックス
掲示板
掲示板
対クライアント装置メッセージ交換手段
サーバ装置間メッセージ交換手段
2N
対クライアント装置メッセージ交換手段
サーバ装置間メッセージ交換手段
2A
3A
3N
20
対電子メールサーバ装置メッセージ交換手段
対電子メールサーバ通信装置
71
72
処理装置
記憶装置
7
インターフェイス部
74
73
ユーザ処理手段
データ交換手段
9
11
8
アプリケーションプログラム
10
クライアント装置

【図2】


サービス受益対象（ユーザまたはアプリケーションプログラム）
対電子メールサーバ装置メッセージ交換手段
電子メールサーバ装置
S1
IDの入力
S2
IDの記憶
S3
パスワード入力
S4
パスワードの記憶
S5
サーバへIDの一斉送信
S6
IDの受信
S7
登録ユーザの確認
S8
登録ユーザか？
YES
S9
サーバの識別コード送信
S10
サービス提供可能なサーバの識別コード受信
S11
サービス提供可能なサーバリスト作成
S12
サーバの情報提供
S13
接続対象のサーバを選択
S14
サービス提供可能なサーバ存在しない
YES
S15
サービス提供不能通知
S16
接続失敗通知受信
NO
S17
選択したサーバへのログイン要求
S18
ID，パスワードの受信
S19
接続可能か？
NO
YES
S20
接続実行
S21
結果通知
S22
接続結果受信
NO
S23
接続に成功したか
YES
S24
サービス開始通知
S25
サービス開始通知受信

【図3】


サービス受益対象（ユーザまたはアプリケーションプログラム）
対電子メールサーバ装置メッセージ交換手段
電子メールサーバ装置
正常時
S31
サービス要求
S32
サービス要求
S33
サービス要求受信
S34
サービス実行
S36
サービス内容受信
サービス内容返送
S35
S38
サービス内容受領
サービス内容転送
S37
障害発生時
S31
サービス要求
S32
サービス要求
S41
サービス内容受信失敗
S12
S11
サービス提供可能なサーバリスト作成
サーバの情報提供
S13
接続対象のサーバを選択
S16a
S15
S14
サーバ障害報告受信
サービス提供不能通知
サービス提供可能なサーバ存在しない
YES
NO
S18
選択したサーバへのログイン要求
ID，パスワードの受信
S19
S17
接続可能か？
NO
YES
S20
接続実行
S22
接続結果受信
結果通知
S21
NO
接続に成功したか
S33
S23
YES
サービス再要求
サービス要求受信
S34
S32
S36
サービス実行
S38
サービス内容受信
サービス内容返送
S35
サービス内容受領
サービス内容転送
S37

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

Bib Data Sheet

CONFIRMATION NO. 3296

SERIAL NUMBER	FILING OR 371(c) DATE / RULE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
11/065,412	02/23/2005	709	2157	2711-0073

APPLICANTS

David A. Farber, Oak View, CA;
Richard E. Greer, Red Lodge, MT;
Andrew D. Swart, Westlake Village, CA;
James A. Balter, Santa Barbara, CA;

**** CONTINUING DATA ***************************

This application is a CON of 09/612,598 07/07/2000 ABN which is a DIV of 09/021,506 02/10/1998 PAT 6,185,598

**** FOREIGN APPLICATIONS *********************

IF REQUIRED, FOREIGN FILING LICENSE GRANTED
** 03/23/2005

		STATE OR COUNTRY	SHEETS DRAWING	TOTAL CLAIMS	INDEPENDENT CLAIMS
Foreign Priority claimed	☐ yes ☐ no	CA	6	1	1
35 USC 119 (a-d) conditions met	☐ yes ☐ no ☐ Met after Allowance				
Verified and Acknowledged	Examiner's Signature Initials				

ADDRESS
42624

TITLE
SHARED CONTENT DELIVERY INFRASTRUCTURE

FILING FEE RECEIVED		
2410	FEES: Authority has been given in Paper No. _____ to charge/credit DEPOSIT ACCOUNT No. _____ for following:	☐ All Fees ☐ 1.16 Fees (Filing) ☐ 1.17 Fees (Processing Ext. of time) ☐ 1.18 Fees (Issue) ☐ Other _____ ☐ Credit

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Approved for use through 7/31/2006. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

PTO/SB/06 (12-04)

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 1V065412

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE ($)	SMALL ENTITY FEE ($)	OR	OTHER THAN SMALL ENTITY RATE ($)	OTHER THAN SMALL ENTITY FEE ($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))							
SEARCH FEE (37 CFR 1.16(k), (l), or (m))							
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))							
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 = *		x 25 =		OR	x 50 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 = *		x 100 =		OR	x 200 =	
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))			180			360	
			TOTAL			TOTAL	

* If the difference in column 1 is less than zero, enter "0" in column 2.

APPLICATION AS AMENDED – PART II

AMENDMENT

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE ($)	SMALL ENTITY ADDITIONAL FEE ($)	OR	OTHER THAN SMALL ENTITY RATE ($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	* 48	Minus	** 20	= 28	x 25 =		OR	x 50 =	$1400
Independent (37 CFR 1.16(h))	* 5	Minus	*** 3	= 2	x 100 =		OR	x 200 =	$400 pd
Application Size Fee (37 CFR 1.16(s))							OR	360	$360 pd
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))					180				
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

AMENDMENT

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)	OR	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
Application Size Fee (37 CFR 1.16(s))							OR		
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))									
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

CLAIMS ONLY

Application Number: 11/065,412

Filing Date:

Applicant(s):

* May be used for additional claims or amendments

CLAIMS	AS FILED		AFTER FIRST AMENDMENT		AFTER SECOND AMENDMENT		*		*		*	
	Indep	Depend	Indep	Depend	Indep	Depend	Indep	Depend	Indep	Depend	Indep	Depend
1												
2	1											
3		1										
4		1										
5		1										
6		1										
7		1										
8		1										
9		1										
10		8										
11		9										
12	1											
13		1										
14		1										
15		1										
16		4										
17		4										
18	1											
19		1										
20		2										
21		2										
22	1											
23		1										
24		2										
25		2										
26	1											
27												
28												
29												
30												
31												
32												
33												
34												
35												
36												
37												
38												
39												
40												
41												
42												
43												
44												
45												
46												
47												
48												
49												
50												
Total Indep	5											
Total Depend	43											
Total Claims	48											

	Indep	Depend	Indep	Depend	Indep	Depend
51						
52						
53						
54						
55						
56						
57						
58						
59						
60						
61						
62						
63						
64						
65						
66						
67						
68						
69						
70						
71						
72						
73						
74						
75						
76						
77						
78						
79						
80						
81						
82						
83						
84						
85						
86						
87						
88						
89						
90						
91						
92						
93						
94						
95						
96						
97						
98						
99						
100						
Total Indep						
Total Depend						
Total Claims						

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:	Atty. Docket No.: 2711-0073
FARBER, David et al.	Confirmation No.: 3296
Appln. No.: 11/065,412	Group Art Unit: 2157
Filed: February 23, 2005	Examiner: Unassigned
Title: SHARED CONTENT DELIVERY INFRASTRUCTURE (as amended)	Date: June 26, 2007

PRELIMINARY AMENDMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Prior to examination on its merits, please amend this application as follows:

Amendments to the title begin on page 2.

Amendments to the Abstract begin on page 3.

Amendments to the claims begin on page 4.

IN THE TITLE

Please replace the title with the following new title:

"Shared Content Delivery Infrastructure"

IN THE ABSTRACT OF THE DISCLOSURE

Please replace the Abstract with the attached new Abstract of the Disclosure.

IN THE CLAIMS

Please cancel claim 1 and add the following new claims 2-22:

1. (Canceled)

2. (NEW) A method, in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some resources via one or more content sources associated with that content provider, the method comprising:

(A) responsive to a first request from a first client for a first resource, said first resource being associated with a first content provider of said plurality of content providers, and said first resource including at least one reference to at least one other resource, said at least one other resource also being associated with said first content provider,

causing said at least one other resource to be served from a shared content delivery network (CDN) formed from a plurality of servers distinct from said content sources associated with said content providers, wherein said first resource is served to the first client from a content source associated with said first content provider, and

(B) responsive to a second request from a second client for a second resource, said second resource being associated with a second content provider of said plurality of content providers, said second content provider being distinct from said first content provider, and said second resource including at least one reference to at least a third resource, said third resource also being associated with said second content provider,

causing the third resource to be served to the second client from the shared content delivery network, wherein the second resource is served to the second client from a content source associated with said second content provider.

3. (NEW) A method as in claim 2 wherein said first resource is a web page associated with said first content provider and wherein said at least one other resource reference in said first resource is an embedded object reference.

4. (NEW) A method as in claim 3 wherein said first request from said first client comprises a first Uniform Resource Locator (URL), and

wherein said second request from said second client comprises a second URL, and

wherein said embedded object reference in said first resource comprises a third URL; and

wherein said reference to said third resource comprises a fourth URL.

5. (NEW) A method as in claim 4

wherein the content source associated with the first content provider is an origin server associated with the first content provider and wherein said first URL includes a first hostname that resolves to said origin server associated with said first content provider, and

wherein the content source associated with the second content provider is an origin server associated with the second content provider, and wherein said second URL includes a second hostname that resolves to said origin server associated with said second content provider, and

wherein said third URL includes a third hostname that resolves to the shared CDN, and wherein said fourth URL includes a fourth hostname that also

resolves to the shared CDN, said third hostname being distinct from said fourth hostname.

6. (NEW) A method as in claim 2 further comprising:

at a first server in said shared CDN:

(a1) if said at least one other resource is not available at said first server, said first server obtaining said at least one other resource; and

(a2) said first server serving said at least one other resource to said first client.

7. (NEW) A method as in claim 6 wherein said first server obtains said at least one other resource from a peer or from at least one content source associated with said first content provider.

8. (NEW) A method as in claim 6 wherein said at least one reference to said at least one other resource comprises a Uniform Resource Locator (URL), and wherein said first server tries to obtain said at least one other resource from a location specified in said URL.

9. (NEW) A method as in claim 8 wherein said location specified in said URL is an content source associated with said first content provider.

10. (NEW) A method as in any one of claims 2-9 wherein said at least one other resource referenced in said first resource is selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources; and

wherein said third resource is a resource selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

11. (NEW) A method as in any one of claims 2-4 and 6-9

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

12. (NEW) A method, in a system in which a plurality of content providers provide multiple resources to multiple clients, and wherein each content provider provides at least some resources via one or more content sources associated with that content provider, and wherein a first content provider of said plurality of content providers has a first web page associated therewith, said first web page including at least a first referenced resource associated with said first content provider; and wherein a second content provider of said plurality of content providers has a second web page, said second web page including at least a second referenced resource associated with said second content provider, the method comprising:

(A) providing a network of repeater servers to form a shared content delivery network (CDN);

(B) responsive to a request for said first web page, after said first web page has been served from a content source associated with said first content provider, serving said first referenced resource from said shared CDN, and

(C) responsive to a request for said second web page, after said second web page has been served from a content source associated with said second content provider, serving said second referenced resource from said shared CDN,

wherein said content sources associated with said first and said second content providers are distinct from said repeater servers.

13. (NEW) A method as in claim 12 wherein said first referenced resource is served from a first server in said shared CDN, the method further comprising:

if said first server does not have a copy of the first referenced resource, then trying to obtain a copy of said first referenced resource.

14. (NEW) A method as in claim 13 wherein said first server in said shared CDN obtains a copy of the first referenced resource from said content source associated with said first content provider.

15. (NEW) A method as in claim 12 wherein

said request for said first web page comprises a first Uniform Resource Locator (URL) having a first hostname; and

said request for said second web page comprises a second URL having a second hostname distinct from said first hostname; and

said first referenced resource is referenced by a third URL having a third hostname distinct from said first hostname; and

said second referenced resource is referenced by a fourth URL having a fourth hostname distinct from said second hostname, and

wherein said first hostname resolves to said content source associated with said first content provider, and

wherein said second hostname resolves to said content source associated with said second content provider, and

wherein said third hostname and said fourth hostname both resolve to said shared CDN.

16. (NEW) A method as in any one of claims 12-15 wherein said first referenced resource is selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources; and

wherein said second referenced resource selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

17. (NEW) A method as in claim 16

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

18. (NEW) An Internet method, operable in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some of its resources via one or more content sources associated with that content provider, wherein a first web page includes a first referenced resource; and wherein a second web page includes a second referenced resource, the method comprising:

(A) providing a network of repeater servers;

(B) after said first web page has been served from a first content source, at a first repeater server in said network of repeater servers:

(b1) if said first repeater server does not have a copy of said first referenced resource, obtaining a copy of said first referenced resource from said first content source; and

(b2) serving said first referenced resource from said first repeater server; and

(C) after said second web page has been served from a second content source, at a second repeater server in said network of repeater servers:

(c1) if said second repeater server does not have a copy of said second referenced resource, obtaining a copy of said second referenced resource from said second content source; and

(c2) serving said second referenced resource from said second repeater server,

wherein said first content source and said content source server are distinct from said network of repeater servers.

19. (NEW) A method as in claim 18 wherein

said first referenced resource is referenced by a first URL having a first hostname; and

said second referenced resource is referenced by a second URL having a second hostname distinct from said first hostname, and

wherein said first hostname and said second hostname both resolve to said network of repeater servers.

20. (NEW) A method as in any one of claims 18-19 wherein said first referenced resource and said second referenced resource are selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

21. (NEW) A method as in claim 20

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

22. (NEW) An Internet method, in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some of its resources via one or more content sources associated with that content provider, wherein a first web page is

identified by a first Uniform Resource Locator (URL) including a first hostname, said first hostname resolving to a first content source; and wherein a second web page is identified by a second URL including a second hostname, said second hostname resolving to a second content source, and wherein the first web page includes a first embedded object, and wherein the second web page includes a second embedded object, the method comprising:

(A) providing a shared network of repeater servers, said first content source and said second content source being distinct from said network of repeater servers;

(B) responsive to a request for said first web page, said request including said first URL, after said first web page has been served via said first content source, at a first repeater server in said network of repeater servers:

(b1) if said first repeater server does not have a copy of said first embedded object, obtaining a copy of said first embedded object from said first content source; and

(b2) serving said first embedded object from said first repeater server; and

(C) after said second web page has been served from said second content source, at a second repeater server in said network of repeater servers:

(c1) if said second repeater server does not have a copy of said second embedded object, obtaining a copy of said second embedded object from said second content source; and

(c2) serving said second embedded object from said second repeater server.

23. (NEW) A method as in claim 22 wherein

the first embedded object is identified in the first web page by a third URL comprising a third hostname, and

the second embedded object is identified in the second web page by a fourth URL comprising a fourth hostname, and

wherein said third hostname and said fourth hostname both resolve to said network of repeater servers.

24. (NEW) A method as in any one of claims 22-23 wherein said first embedded object and said second embedded object are selected from the group comprising static and dynamic: video content, audio content, text, image content.

25. (NEW) A method as in any one of claims 22-23 wherein said

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

26. (NEW) A method for delivering resources to clients in a distributed computing environment, the method comprising:

(A) associating a first referenced resource referenced by a first primary resource with a shared network formed by a plurality of repeater servers, wherein the first primary resource is associated with a first content provider and is maintained on a first content source distinct from the plurality of repeater servers;

(B) associating a second referenced resource referenced by a second primary resource with the shared network, wherein the second primary resource is associated with a second content provider and is maintained on a second content

source distinct from the plurality of repeater servers and distinct from the first content source;

(C) responsive to a first request that causes the first primary resource to be served to a first client from the first content source:

(c1) selecting at least a first repeater server to serve the first referenced resource to the first client;

(c2) responsive to the first repeater server being requested to serve the first referenced resource:

(c21) checking to determine whether a copy of the first referenced resource is available on the first repeater server;

(c22) if a copy of the first referenced resource is available on the first repeater server, attempting to serve the copy of the first referenced resource to the client from the first repeater server; and

(c23) if a copy of the first referenced resource is not available on the first repeater server, replicating the first referenced resource on the first repeater;

(D) responsive to a second request that causes the second primary resource to be served to a second client from the second content source:

(d1) selecting at least a second repeater server to serve the second referenced resource to the second client;

(d2) responsive to the second repeater server being requested to serve the second referenced resource:

(d21) checking to determine whether a copy of the second referenced resource is available on the second repeater server;

(d22) if a copy of the second referenced resource is available on the second repeater server, attempting to serve the copy of the

second referenced resource to the client from the second repeater server; and

(d23) if a copy of the second referenced resource is not available on the second repeater server, replicating the second referenced resource on the second repeater.

REMARKS

By this preliminary amendment, the Title and Abstract of the Disclosure have been changed, claim 1 has been canceled and new claims 2-26 have been added. No new matter has been added.

Claims 2-26 are pending in this application.

An early action and allowance of this application are respectfully requested.

CHARGE STATEMENT: Deposit Account No. 501860, order no. **2711-0073**.
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/ Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached.
This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.

CUSTOMER NUMBER 42624	Respectfully submitted, By: /Brian Siritzky/ Brian Siritzky, Ph.D., Reg. No. 37,497
Davidson Berquist Jackson & Gowdey LLP 4300 Wilson Blvd., 7th Floor, Arlington, Virginia 22203 Main: (703) 894-6400 ● FAX: (703) 894-6430	

ABSTRACT OF THE DISCLOSURE

A system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some resources via one or more content sources associated with that content provider. A client requests a resource from a content provider, where the resources includes references to other resources of the content provider. Referenced resources are served to requesting clients from a shared content delivery network (CDN) formed from a plurality of servers distinct from content sources associated with said content providers. The content sources may be origin servers associated with the respective content providers. The referenced resources may be video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:		Atty. Docket No.:	2711-0073
FARBER, David et al.		Confirmation No.:	3296
Appln. No.:	11/065,412	Group Art Unit:	2157
Filed:	February 23, 2005	Examiner:	Unassigned
Title:	SHARED CONTENT DELIVERY INFRASTRUCTURE (AS AMENDED)	Date:	June 26, 2007

AMENDMENT TRANSMITTAL

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an amendment in the above-identified application and includes the attachments hereto. The signature below is treated as the signature to the attachments in the absence of a signature thereto.

FEE REQUIREMENTS

	Claims remaining after amendment		*Highest number previously Paid For*		*Present Extra*	*Rate: Large/Small Entity*	
1. *Total*	33	minus	20	*	= 13	X $50 / $ 25 =	$650
			* If this number is less than 20, enter "20"				
2. *Independent Claims*	5	minus	3	**	= 2	X $200 / $100 =	$400
			** If this number is less than 3, enter "3"				
3. If amendment enters multiple dependent claim(s) into this application for first time (leave this line blank if this is an reissue application)						$360 / $180 =	$360
4. **Original due date:**							
				Check time period and enter appropriate fee			
5. **Petition is hereby made** to extend the original due date to cover the date this response is filed for which the requisite fee is:				☐ 1 month ☐ 2 months. ☐ 3 months ☐ 4 months ☐ 5 months.		$120 / $60 $450 / $225 $1020 / $510 $1590 / $795 $2160 / $1080	$0
6. ☐ Attached is a Petition/Fee under Rule No.						$	$0
7. Other Fee for							$0
8. **Total Fee:**							$1410

This is a 2 page form, continued on next page . . .

9. ☒ Please charge the total fee on line 8 above to our deposit account below under the stated order number.
Our Deposit Account No.: 501860 Our Order No. (Client-Matter No.): 2711-0073

CHARGE STATEMENT: Deposit Account No. 501860, order no. 2711-0073.
The Commissioner is hereby authorized to charge any fee specifically authorized hereafter, or any missing or insufficient fee(s) filed, or asserted to be filed, or which should have been filed herewith or concerning any paper filed hereafter, and which may be required under Rules 16-18 (missing or insufficiencies only) now or hereafter relative to this application and the resulting Official Document under Rule 20, or credit any overpayment, to our Accounting/Order Nos. shown above, for which purpose a duplicate copy of this sheet is attached **This CHARGE STATEMENT does not authorize charge of the issue fee until/unless an issue fee transmittal sheet is filed.**

CUSTOMER NUMBER 42624	Respectfully submitted, By: /Brian Siritzky/ Brian Siritzky, Ph.D.; Reg. No. 37,497

Davidson Berquist Jackson & Gowdey LLP
4300 Wilson Blvd., 7th Floor, Arlington, VA 22203
Main: (703) 894-6400 ● FAX: (703) 894-6430

Electronic Patent Application Fee Transmittal

Application Number:	11065412
Filing Date:	23-Feb-2005
Title of Invention:	Optimized network resource location
First Named Inventor/Applicant Name:	David A. Farber
Filer:	Brian Siritzky
Attorney Docket Number:	2711-0073

Filed as Large Entity

Utility Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Claims in excess of 20	1202	13	50	650
Independent claims in excess of 3	1201	2	200	400
Multiple dependent claims	1203	1	360	360
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**1410**

Electronic Acknowledgement Receipt	
EFS ID:	1913396
Application Number:	11065412
International Application Number:	
Confirmation Number:	3296
Title of Invention:	Optimized network resource location
First Named Inventor/Applicant Name:	David A. Farber
Customer Number:	42624
Filer:	Brian Siritzky
Filer Authorized By:	
Attorney Docket Number:	2711-0073
Receipt Date:	26-JUN-2007
Filing Date:	23-FEB-2005
Time Stamp:	19:09:07
Application Type:	Utility

Payment information:

Submitted with Payment	yes
Payment was successfully received in RAM	$1410
RAM confirmation Number	3216
Deposit Account	501860

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)	Multi Part /.zip	Pages (if appl.)

1	Preliminary Amendment	2711-0073_Preliminary_Amendment.pdf	323796	no	17
Warnings:					
Information:					
2	Miscellaneous Incoming Letter	Amendment_Transmittal_06-26-2007.pdf	417090	no	2
Warnings:					
Information:					
3	Fee Worksheet (PTO-06)	fee-info.pdf	8383	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	749269		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:	Atty. Docket No.: 2711-0073
FARBER, David et al.	Confirmation No.: 3296
Appln. No.: 11/065,412	Group Art Unit: 2157
Filed: February 23, 2005	Examiner: Unassigned
Title: SHARED CONTENT DELIVERY INFRASTRUCTURE (as amended)	Date: June 26, 2007

PRELIMINARY AMENDMENT

Hon. Commissioner of Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Prior to examination on its merits, please amend this application as follows:

Amendments to the title begin on page 2.

Amendments to the Abstract begin on page 3.

Amendments to the claims begin on page 4.

IN THE TITLE

Please replace the title with the following new title:

"Shared Content Delivery Infrastructure"

IN THE ABSTRACT OF THE DISCLOSURE

Please replace the Abstract with the attached new Abstract of the Disclosure.

IN THE CLAIMS

Please cancel claim 1 and add the following new claims 2-22:

1. (Canceled)

2. (NEW) A method, in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some resources via one or more content sources associated with that content provider, the method comprising:

(A) responsive to a first request from a first client for a first resource, said first resource being associated with a first content provider of said plurality of content providers, and said first resource including at least one reference to at least one other resource, said at least one other resource also being associated with said first content provider,

causing said at least one other resource to be served from a shared content delivery network (CDN) formed from a plurality of servers distinct from said content sources associated with said content providers, wherein said first resource is served to the first client from a content source associated with said first content provider, and

(B) responsive to a second request from a second client for a second resource, said second resource being associated with a second content provider of said plurality of content providers, said second content provider being distinct from said first content provider, and said second resource including at least one reference to at least a third resource, said third resource also being associated with said second content provider,

causing the third resource to be served to the second client from the shared content delivery network, wherein the second resource is served to the second client from a content source associated with said second content provider.

3. (NEW) A method as in claim 2 wherein said first resource is a web page associated with said first content provider and wherein said at least one other resource reference in said first resource is an embedded object reference.

4. (NEW) A method as in claim 3 wherein said first request from said first client comprises a first Uniform Resource Locator (URL), and

wherein said second request from said second client comprises a second URL, and

wherein said embedded object reference in said first resource comprises a third URL; and

wherein said reference to said third resource comprises a fourth URL.

5. (NEW) A method as in claim 4

wherein the content source associated with the first content provider is an origin server associated with the first content provider and wherein said first URL includes a first hostname that resolves to said origin server associated with said first content provider, and

wherein the content source associated with the second content provider is an origin server associated with the second content provider, and wherein said second URL includes a second hostname that resolves to said origin server associated with said second content provider, and

wherein said third URL includes a third hostname that resolves to the shared CDN, and wherein said fourth URL includes a fourth hostname that also

resolves to the shared CDN, said third hostname being distinct from said fourth hostname.

6. (NEW) A method as in claim 2 further comprising:

at a first server in said shared CDN:

(a1) if said at least one other resource is not available at said first server, said first server obtaining said at least one other resource; and

(a2) said first server serving said at least one other resource to said first client.

7. (NEW) A method as in claim 6 wherein said first server obtains said at least one other resource from a peer or from at least one content source associated with said first content provider.

8. (NEW) A method as in claim 6 wherein said at least one reference to said at least one other resource comprises a Uniform Resource Locator (URL), and wherein said first server tries to obtain said at least one other resource from a location specified in said URL.

9. (NEW) A method as in claim 8 wherein said location specified in said URL is an content source associated with said first content provider.

10. (NEW) A method as in any one of claims 2-9 wherein said at least one other resource referenced in said first resource is selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources; and

wherein said third resource is a resource selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

11. (NEW) A method as in any one of claims 2-4 and 6-9

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

12. (NEW) A method, in a system in which a plurality of content providers provide multiple resources to multiple clients, and wherein each content provider provides at least some resources via one or more content sources associated with that content provider, and wherein a first content provider of said plurality of content providers has a first web page associated therewith, said first web page including at least a first referenced resource associated with said first content provider; and wherein a second content provider of said plurality of content providers has a second web page, said second web page including at least a second referenced resource associated with said second content provider, the method comprising:

(A) providing a network of repeater servers to form a shared content delivery network (CDN);

(B) responsive to a request for said first web page, after said first web page has been served from a content source associated with said first content provider, serving said first referenced resource from said shared CDN, and

(C) responsive to a request for said second web page, after said second web page has been served from a content source associated with said second content provider, serving said second referenced resource from said shared CDN,

wherein said content sources associated with said first and said second content providers are distinct from said repeater servers.

13. (NEW) A method as in claim 12 wherein said first referenced resource is served from a first server in said shared CDN, the method further comprising:

if said first server does not have a copy of the first referenced resource, then trying to obtain a copy of said first referenced resource.

14. (NEW) A method as in claim 13 wherein said first server in said shared CDN obtains a copy of the first referenced resource from said content source associated with said first content provider.

15. (NEW) A method as in claim 12 wherein

said request for said first web page comprises a first Uniform Resource Locator (URL) having a first hostname; and

said request for said second web page comprises a second URL having a second hostname distinct from said first hostname; and

said first referenced resource is referenced by a third URL having a third hostname distinct from said first hostname; and

said second referenced resource is referenced by a fourth URL having a fourth hostname distinct from said second hostname, and

wherein said first hostname resolves to said content source associated with said first content provider, and

wherein said second hostname resolves to said content source associated with said second content provider, and

wherein said third hostname and said fourth hostname both resolve to said shared CDN.

16. (NEW) A method as in any one of claims 12-15 wherein said first referenced resource is selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources; and

wherein said second referenced resource selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

17. (NEW) A method as in claim 16

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

18. (NEW) An Internet method, operable in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some of its resources via one or more content sources associated with that content provider, wherein a first web page includes a first referenced resource; and wherein a second web page includes a second referenced resource, the method comprising:

(A) providing a network of repeater servers;

(B) after said first web page has been served from a first content source, at a first repeater server in said network of repeater servers:

(b1) if said first repeater server does not have a copy of said first referenced resource, obtaining a copy of said first referenced resource from said first content source; and

(b2) serving said first referenced resource from said first repeater server; and

(C) after said second web page has been served from a second content source, at a second repeater server in said network of repeater servers:

(c1) if said second repeater server does not have a copy of said second referenced resource, obtaining a copy of said second referenced resource from said second content source; and

(c2) serving said second referenced resource from said second repeater server,

wherein said first content source and said content source server are distinct from said network of repeater servers.

19. (NEW) A method as in claim 18 wherein

said first referenced resource is referenced by a first URL having a first hostname; and

said second referenced resource is referenced by a second URL having a second hostname distinct from said first hostname, and

wherein said first hostname and said second hostname both resolve to said network of repeater servers.

20. (NEW) A method as in any one of claims 18-19 wherein said first referenced resource and said second referenced resource are selected from the group comprising static and dynamic:

video content, audio content, text, image content, web pages, HTML files, XML files, files in a markup language, documents, hypertext documents, data files, and embedded resources.

21. (NEW) A method as in claim 20

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

22. (NEW) An Internet method, in a system in which a plurality of content providers provide multiple resources to multiple clients, wherein each content provider provides at least some of its resources via one or more content sources associated with that content provider, wherein a first web page is

identified by a first Uniform Resource Locator (URL) including a first hostname, said first hostname resolving to a first content source; and wherein a second web page is identified by a second URL including a second hostname, said second hostname resolving to a second content source, and wherein the first web page includes a first embedded object, and wherein the second web page includes a second embedded object, the method comprising:

(A) providing a shared network of repeater servers, said first content source and said second content source being distinct from said network of repeater servers;

(B) responsive to a request for said first web page, said request including said first URL, after said first web page has been served via said first content source, at a first repeater server in said network of repeater servers:

(b1) if said first repeater server does not have a copy of said first embedded object, obtaining a copy of said first embedded object from said first content source; and

(b2) serving said first embedded object from said first repeater server; and

(C) after said second web page has been served from said second content source, at a second repeater server in said network of repeater servers:

(c1) if said second repeater server does not have a copy of said second embedded object, obtaining a copy of said second embedded object from said second content source; and

(c2) serving said second embedded object from said second repeater server.

23. (NEW) A method as in claim 22 wherein

the first embedded object is identified in the first web page by a third URL comprising a third hostname, and

the second embedded object is identified in the second web page by a fourth URL comprising a fourth hostname, and

wherein said third hostname and said fourth hostname both resolve to said network of repeater servers.

24. (NEW) A method as in any one of claims 22-23 wherein said first embedded object and said second embedded object are selected from the group comprising static and dynamic: video content, audio content, text, image content.

25. (NEW) A method as in any one of claims 22-23 wherein said

wherein said content source associated with said first content provider comprises an origin server associated with said first content provider; and

wherein said content source associated with said second content provider comprises an origin server associated with said second content provider.

26. (NEW) A method for delivering resources to clients in a distributed computing environment, the method comprising:

(A) associating a first referenced resource referenced by a first primary resource with a shared network formed by a plurality of repeater servers, wherein the first primary resource is associated with a first content provider and is maintained on a first content source distinct from the plurality of repeater servers;

(B) associating a second referenced resource referenced by a second primary resource with the shared network, wherein the second primary resource is associated with a second content provider and is maintained on a second content