



(11)

EP 2 223 470 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
20.02.2019 Bulletin 2019/08

(51) Int Cl.:
H04L 29/08 (2006.01) **H04L 12/707** (2013.01)
H04L 12/741 (2013.01) **H04L 29/12** (2006.01)

(21) Application number: **08867775.2**

(86) International application number:
PCT/US2008/086704

(22) Date of filing: **12.12.2008**

(87) International publication number:
WO 2009/085670 (09.07.2009 Gazette 2009/28)

(54) DNS WILDCARD BEACONING TO DETERMINE CLIENT LOCATION AND RESOLVER LOAD FOR GLOBAL TRAFFIC LOAD BALANCING

DNS-WILDCARD-BEACONING ZUR BESTIMMUNG EINES KUNDENSTANDORTES UND
RESOLVER FÜR GLOBALEN VERKEHRLASTENAUSGLEICH

BALISAGE DE REMPLACEMENT DNS PERMETTANT DE DÉTERMINER UN EMPLACEMENT DE
CLIENT ET UNE CHARGE DE RÉSOLVEUR POUR ÉQUILIBRAGE DE CHARGE DE TRAFIC
MONDIAL

(84) Designated Contracting States:
**AT BE BG CH CY CZ DE DK EE ES FI FR GB GR
HR HU IE IS IT LI LT LU LV MC MT NL NO PL PT
RO SE SI SK TR**

• **APGAR, David**
San Francisco, California 94117 (US)

(30) Priority: **20.12.2007 US 962051**

(74) Representative: **Boult Wade Tennant LLP**
Verulam Gardens
70 Gray's Inn Road
London WC1X 8BT (GB)

(43) Date of publication of application:
01.09.2010 Bulletin 2010/35

(56) References cited:
US-A1- 2002 038 360 US-A1- 2002 059 622
US-A1- 2002 152 309 US-A1- 2003 065 763
US-A1- 2006 235 972 US-A1- 2007 233 851
US-B1- 6 249 801

(73) Proprietor: **Excalibur IP, LLC**
New York, NY 10017 (US)

(72) Inventors:
• **CHRISTIAN, Michael**
Pacifica, California 94044 (US)

EP 2 223 470 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

FIELD OF THE INVENTION

[0001] The present invention relates to global traffic load balancing in a network.

BACKGROUND

[0002] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

[0003] As used herein, the term "data center" refers to a colocation of associated servers. The servers that belong to a particular data center are within the same building or complex but data centers are typically located geographically distant from each other. The geographic distance adds protection so that catastrophic failure in one data center caused by a natural disaster would not also cause failure in the other data center. For example, one data center might be located on the East Coast in New York and another data center might be located on the West Coast in San Francisco.

[0004] Global load balancing or "GLB," is a mechanism for distributing client access to particular services across a plurality of servers. For example, in a situation in which a particular service is provided by servers that belong to data centers in New York and San Francisco, GLB might distribute client access so that the number of clients connected to the data center in New York is about the same as the number of clients connected to the data center in San Francisco.

[0005] When used in the context of the Internet, GLB may use a variety of active and passive monitoring techniques to generate a complex map of the Internet. Based upon this map, GLB makes traffic routing decisions to connect a client to the "closest" server. As used herein, "close" does not necessarily mean basing the determination only on geographic proximity. As used herein, a "close" server is a server that results in the fastest connection to the client. Thus, if a server that was located 100 miles away were slower for the client to reach than a server located 200 miles away because of heavy congestion, then GLB would route the client to the "closer" server that is 200 miles away.

[0006] If a user wishes to connect to a web application or a web page, a DNS request is made. A DNS request begins with a user making a request through a client machine, often by typing a domain (e.g. "www.sampledomain.com") in a web browser. The request is sent from the client to a service provider's local DNS resolver ("LDNS"). An LDNS resolver accepts the request from the client and responds to the request with the domain's IP address if the LDNS resolver has stored the answer

to the request in a cache. If the LDNS does not have the answer stored, the LDNS forwards the request to an authoritative DNS resolver. An authoritative DNS resolver is a server that maintains data for the network of a domain. The authoritative DNS resolver receives requests from a LDNS resolver and replies to the LDNS resolver with an IP address of a particular server to connect with the domain. As used herein, GLB resolvers may be a part of an authoritative DNS resolver. GLB resolvers may also be separate or located remotely from an authoritative name server. This may vary from implementation to implementation.

[0007] Unfortunately, GLB decisions may be based upon insufficient data. For example, if a request is made from the LDNS resolver to the GLB, the request contains information only about the LDNS and not the client that originated the request. Thus, the GLB is forced to make a decision based upon the location and congestion at the LDNS resolver rather than at the client.

[0008] Basing routing information on the LDNS resolver and not the client may cause two major problems. First, the client may be located very differently, by geography and network, than the LDNS resolver. This often leads to incorrect proximity mapping by the GLB. Second, because the LDNS resolver caches replies, the GLB is unable to determine the quantity of requests that are being generated by clients sitting behind the LDNS resolver. A single user performing a DNS lookup and one million different users may generate the same amount of traffic at the GLB. This makes load balancing determinations very inaccurate.

[0009] An example of GLB based upon DNS is illustrated in Fig. 1. In Fig. 1, two data centers, or colocations, are located in geographically distinct areas. One data center is located in New York 103 and the other data center is located in San Francisco 101. The data center in New York 103 has an IP address 1.2.3.4 and the data center in San Francisco 101 has an IP address of 2.2.2.2.

[0010] A client 105 wishes to connect to the web page, "www.sampleconnection.com," hosted by the two data centers. The client 105, which uses ACME internet service provider, makes a request to connect to the domain. The request from the client 105 is sent 111 first to the ACME LDNS resolver 107. Based upon the request, the LDNS may have the answer stored in cache or forward the request to the domain's authoritative server that provides the IP address of the domain.

[0011] The ACME LDNS resolver 107 checks whether the IP address of the domain is stored in the LDNS resolver's cache. If the IP address is stored in the cache, then the stored IP address is sent to the client in response to the request. If the IP address to the domain is not found in the cache, then the LDNS resolver sends a request to the authoritative DNS resolver for the domain "www.sampleconnection.com" to obtain an IP address. This is shown in the request 115 made to the GLB and authoritative DNS resolver 109 for the domain "www.sampleconnection.com."

[0012] The GLB (with the authoritative DNS resolver) 109 then determines, based upon the request for the web page, whether to return the IP address of 1.2.3.4 for New York or 2.2.2.2 for San Francisco. Logic in the GLB analyzes the request including the location where the request originated, and the availability and load of the servers to determine which particular server is best for a particular client.

[0013] Unfortunately, the authoritative DNS server 109 views the request as originating from the location of the ACME LDNS resolver 107 in North Carolina, and not from the location of the client 105 in Arizona. Based upon the information from the LDNS resolver 107, the authoritative DNS server 109 might select a connection server based upon a request from North Carolina and the much geographically longer path to New York 119 might be selected rather than the shorter path from the client 105 to San Francisco 117.

[0014] The authoritative DNS resolver 109 is also unable to determine the number of clients that may reside behind a LDNS resolver. A single request to the authoritative DNS resolver may actually be for many clients behind the LDNS resolver. For example, the ACME LDNS 107 might make a single request for client 105 for the domain. After the first request, the ACME LDNS 107 stores the IP address of the domain in cache. One second later, ten thousand more requests are made for the same domain. Because the IP address is stored in the LDNS resolver's cache, the IP address is automatically returned to the clients and no additional requests are sent to the authoritative DNS server, for a period of time specified by the Time To Live setting in the response. Thus, should a server become overloaded or fail and the authoritative DNS server and GLB 109 must transfer clients from the overloaded or failed server to a healthy server, the authoritative DNS server 109 is unable to determine the number of clients sitting behind the LDNS resolver 107 and proper load balancing may not be maintained.

[0015] In US6249801 there is disclosed a method for load balancing requests on a network.

SUMMARY OF INVENTION

[0016] The invention is defined by the independent claims. Various aspects and embodiments of the invention are set out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a diagram displaying a client connecting to a web application via DNS-based GLB to servers in different data centers;

FIG. 2 is a diagram displaying how an aggregation

server maps IP addresses of a client to the IP address of an LDNS resolver based upon web beacons, according to an embodiment of the invention; FIG. 3 is a diagram displaying connection time calculation tables based upon IP addresses of clients and particular colocations, according to an embodiment of the invention; and

FIG. 4 is a block diagram of a computer system on which embodiments of the invention may be implemented.

DETAILED DESCRIPTION

[0018] Techniques are described to determine the number and location of clients that are residing behind an LDNS resolver and to use that information in traffic routing decisions. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

GENERAL OVERVIEW

[0019] In an embodiment, DNS wildcard beaconing is used in order to generate data to determine the number and location of clients that reside behind particular LDNS resolvers. The data is then used to estimate the load that is being generated by a particular LDNS resolver and where the particular LDNS resolver's clients are located in the network. This information allows the GLB server to make more educated and accurate routing decisions.

DNS WILDCARD BEACONING

[0020] In an embodiment, a DNS wildcard beacon is initiated from an existing web application. As used herein, web beacons are also known as pixel tags, clear GIFs, or zero-content images and are transparent or invisible graphic images, usually no larger than a 1 x 1 pixel. In an embodiment, a web beacon is generated by placing a small amount of code, or beacon code, into production web pages. Clients request and are served the production web page from a web server in a data center. When the production web page is processed by the client, the beacon code causes the client browser to retrieve web beacons from beacon servers specified in the beacon code. The requests are performed in the background of the client so no interference occurs with the current web page load from the web server.

[0021] In an embodiment, a globally unique wildcard hostname is generated for a web beacon download and that beacon URL is embedded in an existing web result and sent to the client which proceeds to download that image. In another embodiment, a globally unique wild-

card hostname is generated for a HTTP 204 (i.e., "no content") download and placed in a web result. For example, the following globally unique wildcard hostname might be generated:

"http://12093898978.dnsb.company.com/onepixel.gif"

The globally unique wildcard hostname comprises (1) a unique alphanumeric sequence ("12093898978"), (2) a domain name ("*.dnsb.company.com"), and (3) a name of the web beacon object to be retrieved ("onepixel.gif").

[0022] When a client begins the process to download the beacon image, the client must first resolve the unique hostname into an IP address. The client sends a request to the client's LDNS resolver. The LDNS resolver examines the LDNS resolver's cache to determine whether an IP address is available for the hostname requested in the request. If the IP address is available, the LDNS resolver would respond to the request with an IP address of the hostname. However, since the hostnames for the web beacons are globally unique, the IP address should not be available in the LDNS resolver's cache.

[0023] Because the IP address of the hostname is unavailable from the cache, the LDNS resolver sends a request to the authoritative DNS resolver. In an embodiment, the authoritative DNS resolver is also a beacon resolver for the domain requested (in this case, "dnsb.company.com"). As used herein, a "beacon resolver" is a server that measures and logs statistics based upon the requests received at the authoritative DNS resolver. The beacon resolver also returns an IP address to which the client may retrieve the web beacon.

[0024] In an embodiment, the authoritative DNS resolvers/ beacon resolvers are configured to respond to requests for the domain ("*.dnsb.company.com") with the IP address of a beacon collection server. As this is performed, the beacon resolvers log the unique wildcard hostname (e.g., "12093898978.dnsb.company.com") and LDNS resolvers' IP addresses for all such requests. In another embodiment, the beacon resolver records the unique alphanumeric sequence ("12093898978") located within the hostname and the LDNS resolver's IP address. The LDNS resolver/ beacon resolver then sends the IP address of the beacon collection server to the client.

[0025] The client then submits a request to the beacon collection server for the requested URL, with the request including the globally unique hostname. The request by the client may be any Transmission Control Protocol ("TCP")-based protocol, including, but not limited to, HTTP, FTP, or any other communications protocol that is based on TCP. In an embodiment, the client submits an HTTP request to the beacon collection server for the requested URL, inserting the same globally unique hostname (i.e., "12093898978.dnsb.company.com") in the HTTP host header. The beacon collection server receives the HTTP request from the client. The beacon collection server logs the unique wildcard hostname and the IP address of the client. In another embodiment, the beacon collection server logs the unique alphanumeric

sequence in the hostname and the IP address of the client. The beacon server responds to the request by sending a single pixel, zero-content image, or HTTP 204 to the client, depending upon the nature of the request. Additional connection quality statistics may be logged and measured by a kernel module within the beacon resolver.

[0026] A diagram illustrating web beacons and how web beacons are used to map clients to a particular LDNS resolver, according to an embodiment, is shown in Fig. 2. The process begins when client 200 requests a web page and is served the web page from web server 202. Within the web page is beacon code, that instructs client 200 to retrieve a web beacon, or a zero-content image, from a location specified in the beacon code. When the web page is processed by the client 200, retrieval of the web beacon begins.

[0027] A globally unique hostname (e.g., "12345.dnsb.company.com") is used as the location of the web beacon. Client 200 sends the request for the web beacon to LDNS resolver 204, that is owned by the ISP of client 200, in order to resolve the hostname into an IP address. LDNS resolver first checks if the IP address is available within the LDNS cache. As the hostname is unique, the LDNS resolver forwards the request to the authoritative DNS resolver for the hostname. The authoritative DNS resolver is also the beacon resolver 206 for the hostname (i.e., "*.dnsb.company.com"). The beacon resolver 206 logs the IP address of the LDNS resolver 204 and the unique wildcard of the hostname (i.e., "12345"). The beacon resolver 206 then returns to the LDNS resolver 204 the IP address of the beacon collection server 208 where the client may request the web beacon. The LDNS resolver 204 forwards the IP address of the beacon collection server 208 to client 200. Client 200 then sends the HTTP request ("12345.dnsb.company.com") to the beacon collection server 208. The beacon collection server 208 receives the request from the client and logs the IP address of client 200 and the unique wildcard of the hostname (i.e., "12345"). The beacon collection server 208 responds to the client 200 with a zero-content image or an HTTP 204 (no download) status.

AGGREGATING AND PROCESSING DATA

[0028] In an embodiment, an aggregation server collects data from the beacon resolver and the beacon collection server. The data from the beacon resolver contains a globally unique hostname and the IP address of an LDNS. The data from the beacon collection server contains a globally unique hostname and the IP address of a client. Using the globally unique hostname as a key, the aggregation server may map the IP addresses of clients to the IP address of a particular LDNS. Thus, by aggregating data over time, the number and IP addresses of clients that reside behind a particular LDNS may be determined. The proximity of a client may also be determined within a network based upon the IP addresses stored for the clients. This data may then be exported

from the aggregation server to GLB servers in order to perform more accurate routing.

[0029] An illustration of this process may be seen in Fig. 2. In Fig. 2, beacon resolver 206 has logged the IP address of the LDNS resolver 204 and the unique sequence in the hostname ("12345") from the request. Beacon collection server has logged the IP address of client 200 and the unique sequence in the hostname ("12345") from the request. Aggregation server 210 collects the logged data from beacon resolver 206 and beacon collection server 208. At the aggregation server 210, the IP address of client 200 is mapped to the IP address of the LDNS resolver 204 using the unique hostname sequence as a key. Thus, in the example, LDNS resolver 204 is associated with unique hostname "12345" and client 200 is associated with unique hostname "12345." Because the hostname ("12345") matches, the IP address of client 200 is mapped to the IP address of LDNS resolver 204. The information that client 200 resides behind LDNS resolver 204 is sent from the aggregation server 210 to the GLB server 212 in order for the GLB server to better route subsequent requests based upon client proximity and load.

[0030] In an embodiment, load on a particular LDNS is accurately determined by finding the number of client IP addresses associated with a given LDNS resolver. For example, aggregated data may show that one particular LDNS resolver is associated with twenty different client IP addresses while another particular LDNS resolver is associated with fifty thousand different client IP addresses. This technique allows accurate determinations of load from an LDNS resolver rather than being forced to make assumptions about the load based on the type and breadth of requests seen from an LDNS resolver.

[0031] In an embodiment, client proximity mapping is improved by analyzing the IP addresses of the clients behind a particular LDNS resolver. Based upon the IP addresses of clients, an approximation of a client's location, geographically and within a network, may be determined. Thus, upon receiving a request from a LDNS resolver, a GLB may base routing decisions upon the locations of clients and not the location of the LDNS. The improvement in routing is greatest when the LDNS resolver is located distantly from the location of the LDNS resolver's associated clients.

LOAD BALANCING SERVER USING AGGREGATION SERVER DATA

[0032] An example follows of how a global load balancer may use the data from the aggregation server. Global load balancing is performed when a client requests to visit a particular web page. A DNS lookup is performed to translate the website URL entered by the client (e.g. "www.yahoo.com") into an IP address (e.g. "209.131.36.158"). The lookup is directed first to the LDNS resolver, and if the LDNS resolver does not possess the information, to an authoritative name server

(that is also a load balancing server). The load balancing server examines the requesting IP address of the LDNS resolver. The requesting IP address is then compared to the information for that particular netblock, or range of IP addresses. The load balancing server selects the first available web server on the list sorted by various statistics such as proximity and load and returns the IP address of the web server to the LDNS resolver. By having additional information about clients behind the LDNS resolver, the load balancing server better routes the original client to the web server with the best connectivity.

[0033] An illustration of how data from the aggregation server is used, according to an embodiment of the invention, is shown in Fig. 3. The data has "IP Address" column 300, "Colocation A" column 302, "Colocation B" column 304, and "Colocation C" column 306. "IP Address" column 300 lists the IP addresses to which each of the colocations connect. In row 308, the IP Address "1.1.1.0" indicates that machines at the IP address "1.1.1.x" where x can be any number between 0-255 may connect to colocation A in 10 ms, to colocation B in 20 ms, and to colocation C in 40 ms.

[0034] In row 310, the IP Address "2.2.2.0" indicates that machines at the IP address "2.2.2.x" where x can be any number between 0-255 may connect to colocation A in 50 ms, to colocation B in 80 ms, and to colocation C in 30 ms.

[0035] The data from the aggregation server might indicate that a LDNS resolver with the IP address "5.5.5.5" comprises clients, half with IP addresses "1.1.1.x" and the other half with IP addresses "2.2.2.x." Under this circumstance, the connection times to the three different colocation centers may be determined by finding the average of the connection times.

[0036] For example, row 312 displays the IP address "5.5.5.0." This row would indicate the connection times for the LDNS resolver. Thus the connection time to colocation A from the LDNS resolver is 30 ms, the average of 10 ms (from "1.1.1.0") and 50 ms (from "2.2.2.0"). The connection time to colocation B from the LDNS resolver is 50 ms, the average of 20 ms (from "1.1.1.0") and 80 ms (from "2.2.2.0"). The connection time to colocation C from the LDNS resolver is 35 ms, the average of 40 ms (from "1.1.1.0") and 30 ms (from "2.2.2.0"). Though the connection times from the LDNS resolver to each colocation center is not exact, by taking into account connection times of clients based upon IP addresses, an accurate estimation may be made.

HARDWARE OVERVIEW

[0037] Figure 4 is a block diagram that illustrates a computer system 400 upon which an embodiment of the invention may be implemented. Computer system 400 includes a bus 402 or other communication mechanism for communicating information, and a processor 404 coupled with bus 402 for processing information. Computer system 400 also includes a main memory 406, such as

a random access memory (RAM) or other dynamic storage device, coupled to bus 402 for storing information and instructions to be executed by processor 404. Main memory 406 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 404. Computer system 400 further includes a read only memory (ROM) 408 or other static storage device coupled to bus 402 for storing static information and instructions for processor 404. A storage device 410, such as a magnetic disk or optical disk, is provided and coupled to bus 402 for storing information and instructions.

[0038] Computer system 400 may be coupled via bus 402 to a display 412, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 414, including alphanumeric and other keys, is coupled to bus 402 for communicating information and command selections to processor 404. Another type of user input device is cursor control 416, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 404 and for controlling cursor movement on display 412. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0039] The invention is related to the use of computer system 400 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 400 in response to processor 404 executing one or more sequences of one or more instructions contained in main memory 406. Such instructions may be read into main memory 406 from another machine-readable medium, such as storage device 410. Execution of the sequences of instructions contained in main memory 406 causes processor 404 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0040] The term "machine-readable medium" as used herein refers to any medium that participates in providing data that causes a machine to operation in a specific fashion. In an embodiment implemented using computer system 400, various machine-readable media are involved, for example, in providing instructions to processor 404 for execution. Such a medium may take many forms, including but not limited to storage media and transmission media. Storage media includes both non-volatile media and volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 410. Volatile media includes dynamic memory, such as main memory 406. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 402. Transmission media can also take the form of acoustic or light waves,

such as those generated during radio-wave and infra-red data communications. All such media must be tangible to enable the instructions carried by the media to be detected by a physical mechanism that reads the instructions into a machine.

[0041] Common forms of machine-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0042] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 404 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 400 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 402. Bus 402 carries the data to main memory 406, from which processor 404 retrieves and executes the instructions. The instructions received by main memory 406 may optionally be stored on storage device 410 either before or after execution by processor 404.

[0043] Computer system 400 also includes a communication interface 418 coupled to bus 402. Communication interface 418 provides a two-way data communication coupling to a network link 420 that is connected to a local network 422. For example, communication interface 418 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 418 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 418 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0044] Network link 420 typically provides data communication through one or more networks to other data devices. For example, network link 420 may provide a connection through local network 422 to a host computer 424 or to data equipment operated by an Internet Service Provider (ISP) 426. ISP 426 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 428. Local network 422 and Internet 428 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the vari-

ous networks and the signals on network link 420 and through communication interface 418, which carry the digital data to and from computer system 400, are exemplary forms of carrier waves transporting the information.

[0045] Computer system 400 can send messages and receive data, including program code, through the network(s), network link 420 and communication interface 418. In the Internet example, a server 430 might transmit a requested code for an application program through Internet 428, ISP 426, local network 422 and communication interface 418.

[0046] The received code may be executed by processor 404 as it is received, and/or stored in storage device 410, or other non-volatile storage for later execution. In this manner, computer system 400 may obtain application code in the form of a carrier wave.

[0047] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

Claims

1. A computer-implemented method, comprising:

receiving a plurality of first data from a first server (206);

wherein each first data of the plurality of first data includes a first unique hostname and data that identifies a particular Local Domain Name Service server;

receiving a plurality of second data from a second server (208) that is not the first server (206); wherein each second data of the plurality of second data includes a second unique hostname and data that identifies a particular client;

wherein, for each first data of the plurality of first data, the first unique hostname included in the first data was provided in a request from a client to the particular Local Domain Name Service server identified in the first data;

wherein, for each second data of the plurality of second data, the second unique hostname included in the second data was provided in a request from the client identified in the second data

to a particular server;
and

aggregating clients by the particular Local Domain Name Service servers they are associated with including, for each second data of the plurality of second data, where the unique hostname of the second data matches the unique hostname of a first data of the plurality of first data, creating and storing mapping data associating the particular client identified in the second data with the particular Local Domain Name Service server identified in the first data of the plurality of first data.

2. The method of Claim 1, further comprising:

outputting each mapping data to one or more load balancing servers for use by said one or more load balancing servers in making routing decisions.

3. The method of Claim 1, wherein the data that identifies the particular Local Domain Name Service server comprises an Internet Protocol address.

4. The method of Claim 1, wherein the data that identifies a particular client comprises an Internet Protocol address.

5. The method of Claim 1, wherein, for each first data of the plurality of first data, the first unique hostname included in the first data was provided in a Domain Name Service request from the client to the particular Local Domain Name Service server identified in the first data.

6. The method of Claim 1, wherein, for each first data of the plurality of first data, the request from the client to the particular Local Domain Name Service server identified in the first data was caused by web beacon code served to the client.

7. The method of Claim 1, wherein, for each second data of the plurality of second data, the second unique hostname included in the second data was provided in a HyperText Transfer Protocol request from the client identified in the second data to the particular server.

8. The method of Claim 1, wherein, for each second data of the plurality of second data, the request from the client identified in the second data to the particular server was caused by web beacon code served to the client.

9. The method of Claim 2, wherein making routing decisions is based upon proximity of clients residing behind Local Domain Name Service servers.

10. The method of Claim 2, wherein making routing decisions is based upon a number of clients residing behind a particular Local Domain Name Service server.

11. The method of Claim 1, wherein, for each second data of the plurality of second data, the second unique hostname included in the second data was provided in a HyperText Transfer Protocol request for a web beacon from the client identified in the second data to the particular server.

12. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to carry out a method according to any one of the preceding claims.

Patentansprüche

1. Computerimplementiertes Verfahren, umfassend: Empfang einer Vielzahl von ersten Daten von einem ersten Server (206); wobei jeweils die ersten Daten der Vielzahl von ersten Daten einen ersten eindeutigen Hostnamen und Daten, die einen bestimmten Local-Domain-Name-Service-Server identifizieren, einschließen; Empfangen einer Vielzahl von zweiten Daten von einem zweiten Server (208), der nicht der erste Server (206) ist; wobei jeweils die zweiten Daten aus der Vielzahl der zweiten Daten einen zweiten eindeutigen Hostnamen und Daten, die einen bestimmten Client identifizieren, einschließen; wobei jeweils für die erste Daten der Vielzahl von ersten Daten der erste eindeutige Hostname, der in den ersten Daten mitumfasst ist, in einer Anforderung von einem Client an den jeweiligen in den ersten Daten identifizierten Local-Domain-Name-Service-Server bereitgestellt wurde; wobei jeweils für die zweiten Daten der Vielzahl von zweiten Daten der zweite eindeutige Hostname, der in den zweiten Daten enthalten ist, in einer Anforderung von dem mit den zweiten Daten identifizierten Client an einen bestimmten Server bereitgestellt wurde; und Aggregation von Clients durch die jeweiligen Local-Domain-Name-Service-Server, denen sie zugeordnet sind, einschließlich, jeweils für zweite Daten der Vielzahl von zweiten Daten, wo der eindeutige Hostname der zweiten Daten jeweils mit dem eindeutigen Hostnamen erster Daten der Vielzahl von ersten Daten übereinstimmt, Erzeugen und Speichern von Mapping-Daten, die den in den zweiten Daten identifizierten bestimmten Client mit dem in den ersten Daten der Vielzahl von ersten Daten identifizierten

speziellen Local-Domain-Name-Service-Server in Verbindung bringen.

2. Verfahren nach Anspruch 1, ferner umfassend:

Ausgeben von Mapping-Daten jeweils an einen oder mehrere Lastausgleichsserver zur Verwendung durch den einen oder die mehreren Lastausgleichsserver bei der Entscheidungsfindung für das Routing.

3. Verfahren nach Anspruch 1, wobei die Daten, die den jeweiligen Local-Domain-Name-Service-Server identifizieren, eine Internet-Protokoll-Adresse umfassen.

4. Verfahren nach Anspruch 1, wobei die Daten, die einen bestimmten Kunden identifizieren, eine Internet-Protokoll-Adresse umfassen.

5. Verfahren nach Anspruch 1, wobei für erste Daten jeweils aus der Vielzahl der ersten Daten der erste eindeutige Hostname, der in den ersten Daten enthalten ist, in einer Domain-Name-Service-Anforderung von dem Client an den jeweiligen Local-Domain-Name-Service-Server, der in den ersten Daten identifiziert ist, bereitgestellt wurde.

6. Verfahren nach Anspruch 1, wobei für erste Daten jeweils aus der Vielzahl der ersten Daten die Anforderung von dem Client an den jeweiligen Local-Domain-Name-Service-Server, der in den ersten Daten identifiziert wurde, durch einen Web-Beacon-Code verursacht wurde, der dem Client bereitgestellt wurde.

7. Verfahren nach Anspruch 1, wobei jeweils für zweite Daten der Vielzahl von zweiten Daten der zweite eindeutige Hostname, der in den zweiten Daten enthalten ist, in einer Hypertext-Übertragungsprotokoll-Anforderung von dem in den zweiten Daten identifizierten Client an den jeweiligen Server bereitgestellt wurde.

8. Verfahren nach Anspruch 1, wobei jeweils für zweite Daten aus der Vielzahl von zweiten Daten die Anforderung von dem in den zweiten Daten identifizierten Client an den jeweiligen Server durch einen dem Client bereitgestellten Web-Beacon-Code bewirkt wurde.

9. Verfahren nach Anspruch 2, wobei das Treffen von Routing-Entscheidungen auf der Nähe von Clients basiert, die sich hinter Local-Domain-Name-Service-Servern befinden.

10. Verfahren nach Anspruch 2, wobei das Treffen von

Routing-Entscheidungen auf einer Anzahl von Clients basiert, die sich hinter einem bestimmten Local-Domain-Name-Service-Server befinden.

11. Verfahren nach Anspruch 1, wobei jeweils für zweite 5
Daten aus der Vielzahl von zweiten Daten der zweite
eindeutige Hostname, der in den zweiten Daten ent-
halten ist, in einer Hypertext-Übertragungsprotokol-
lanforderung für ein Web-Beacon von dem in den
zweiten Daten identifizierten Client an den jeweiligen
Server bereitgestellt wurde. 10
12. Computerlesbares Medium, das eine oder mehrere
Abfolgen von Anweisungen trägt, die, wenn sie von 15
einem oder mehreren Prozessoren ausgeführt wer-
den, bewirken, dass der eine oder die mehreren Pro-
zessoren ein Verfahren nach einem der vorherge-
henden Ansprüche durchführen. 15

Revendications

1. Procédé implémenté par ordinateur, comprenant les 20
étapes consistant à :

recevoir une pluralité de premières données à
partir d'un premier serveur (206) ;
dans lequel chaque première donnée de la plu-
ralité de premières données comprend un pre-
mier nom d'hôte unique et des données qui iden-
tifier un serveur de service de nom de domaine
local particulier ;
recevoir une pluralité de secondes données à
partir d'un second serveur (208) qui n'est pas le
premier serveur (206) ;
dans lequel chaque seconde donnée de la plu-
ralité de secondes données comprend un se-
cond nom d'hôte unique et des données qui
identifient un client particulier ;
dans lequel, pour chaque première donnée de
la pluralité de premières données, le premier
nom d'hôte unique inclus dans les premières
données a été fourni dans une requête prove-
nant d'un client au serveur de service de nom
de domaine local particulier identifié dans les
premières données ;
dans lequel, pour chaque seconde donnée de
la pluralité de secondes données, le second
nom d'hôte unique inclus dans les secondes
données a été fourni dans une requête prove-
nant du client identifié dans les secondes don-
nées à un serveur particulier ; et
agrérer des clients par les serveurs de service
de nom de domaine local particuliers auxquels
ils sont associés, y compris, pour chaque se-
conde donnée de la pluralité de secondes don-
nées, où le nom d'hôte unique de la seconde
donnée correspond au nom d'hôte unique d'une 55

première donnée de la pluralité de premières
données, créer et stocker des données de car-
tographie associant le client particulier identifié
dans les secondes données et le serveur de ser-
vice de nom de domaine local particulier identifié
dans les premières données de la pluralité de
premières données.

2. Procédé selon la revendication 1, comprenant en 10
outre :

la délivrance en sortie de chaque donnée de car-
tographie à un ou plusieurs serveurs d'équilibra-
ge de charge pour une utilisation par lesdits un
ou plusieurs serveurs d'équilibrage de charge
lors de la prise de décisions de routage.

3. Procédé selon la revendication 1, dans lequel les 20
données qui identifient le serveur de service de nom
de domaine local particulier comprennent une adres-
se de protocole Internet.
4. Procédé selon la revendication 1, dans lequel les
données qui identifient un client particulier compren-
nent une adresse de protocole Internet. 25
5. Procédé selon la revendication 1, dans lequel pour
chaque première donnée de la pluralité de premières
données, le premier nom d'hôte unique inclus dans
les premières données a été fourni dans une requête
de service de nom de domaine du client au serveur
de service de nom de domaine local particulier iden-
tifié dans la première donnée. 30
6. Procédé selon la revendication 1, dans lequel pour
chaque première donnée de la pluralité de premières
données, la requête du client au serveur de service
de nom de domaine local particulier identifié dans la
première donnée a été provoquée par un code de
balise Web transmis au client. 35
7. Procédé selon la revendication 1, dans lequel pour
chaque seconde donnée de la pluralité de secondes
données, le second nom d'hôte unique inclus dans
la seconde donnée a été fourni dans une requête de
protocole de transfert HyperText du client identifié
dans la seconde donnée au serveur particulier. 40
8. Procédé selon la revendication 1, dans lequel pour
chaque seconde donnée de la pluralité de secondes
données, la requête provenant du client identifié
dans la seconde donnée au serveur particulier a été
provoquée par un code de balise Web transmis au
client. 45
9. Procédé selon la revendication 2, dans lequel la pri-
se de décisions de routage est basée sur la proximité
de clients résidant derrière des serveurs de service 50

de nom de domaine local.

- 10.** Procédé selon la revendication 2, dans lequel la prise de décisions de routage est basée sur un nombre de clients résidant derrière un serveur de service de nom de domaine local particulier. 5
- 11.** Procédé selon la revendication 1, dans lequel pour chaque seconde donnée de la pluralité de secondes données, le second nom d'hôte unique inclus dans la seconde donnée a été fourni dans une requête de protocole de transfert HyperText pour une balise Web provenant du client identifié dans la seconde donnée au serveur particulier. 10 15
- 12.** Support lisible par ordinateur portant une ou plusieurs séquences d'instructions qui, lorsqu'elles sont exécutées par un ou plusieurs processeurs, amène les un ou plusieurs processeurs à mettre en oeuvre un procédé selon l'une quelconque des revendications précédentes. 20

25

30

35

40

45

50

55

FIG. 1
(Prior Art)




FIG. 2





FIG. 3

<u>300</u>	<u>302</u>	<u>304</u>	<u>306</u>
IP Address	Colocation A	Colocation B	Colocation C
<u>308</u> 1.1.1.0	10 ms	20 ms	40 ms
<u>310</u> 2.2.2.0	50 ms	80 ms	30 ms
<u>312</u> 5.5.5.0	30 ms	50 ms	35 ms

FIG. 4



EP 2 223 470 B1

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 6249801 B [0015]