

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 July 2006 (06.07.2006)

PCT

(10) International Publication Number
WO 2006/072114 A2

- (51) International Patent Classification:
G06F 15/16 (2006.01) G06F 15/173 (2006.01)
- (21) International Application Number:
PCT/US2006/000170
- (22) International Filing Date: 3 January 2006 (03.01.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/027,275 31 December 2004 (31.12.2004) US
11/271,941 10 November 2005 (10.11.2005) US
- (71) Applicant (for all designated States except US): LEVEL
3 COMMUNICATIONS, INC. [US/US]; 1025 Eldorado
Boulevard, Broomfield, Colorado 80021 (US).


- (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,
KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV,
LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI,
NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG,
SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US,
UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

- (72) Inventor; and
- (75) Inventor/Applicant (for US only): SIRKIN, Craig
[US/US]; 3186 Elmira Court, Denver, Colorado 80238
(US).
- (74) Agents: RIETH, Damon, A. et al.; Faegre & Benson,
LLP, 2200 Wells Fargo Center, 90 South Seventh Street,
Minneapolis, Minnesota 55402 (US).

Published:
— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: GEO-LOCATING LOAD BALANCING



(57) Abstract: Methods and apparatus are provided for geo-locating load balancing. According to one embodiment, a communication network architecture includes multiple servers, multiple load balancers, and multiple geographically dispersed communication devices. The servers provide services to the communication devices within the communication network. The load balancers each service a shared virtual Internet Protocol (IP) address common to all of the load balancers and perform load balancing of service requests on behalf of two or more of the servers that are located geographically proximate to the load balancer. The communication devices are communicatively coupled with the load balancers and are configured to issue service requests intended for any of the servers to the shared virtual IP address, whereby, upon issuing a service request, a communication device is directed to a particular server selected by a load balancing routine that is associated with a load balancer that is closest to the communication device.

WO 2006/072114 A2

GEO-LOCATING LOAD BALANCING

BACKGROUND

Field

[0001] Embodiments of the present invention generally relate to the field of load balancing. More particularly, embodiments of the present invention relate to techniques for directing geographically dispersed clients to the closest application servers and balancing the load among such application servers.

Description of the Related Art

[0002] Load balancing generally refers to an attempt to distribute processing and/or communications activity evenly across a computer network so that no single device is overwhelmed. Examples of existing load balancing methodologies include Round-robin Domain Name System (DNS), flow-based load balancing, and Anycast addressing. Round-robin DNS and flow-based load balancing are limited in that they do not factor into the load balancing the location of the client or the host. Meanwhile, Anycast addressing balances only based on network metrics, has scalability issues, and does not ensure that traffic continues to reach the same destination from message to subsequent message.

[0003] **FIG. 1** conceptually illustrates an existing DNS Round-robin approach for finding an available server offering a desired service. In this simplified illustration, clients 101-105 are communicatively coupled with servers 131-135 via a network 120. When a client, such as one of clients 101-105, needs access to a particular service offered by one of servers 131-135, it issues a Hyper Text Transport Protocol (HTTP) request containing a domain name 111 associated with the desired service. A DNS server 110 translates the domain name 111 into a corresponding set of Internet Protocol (IP) addresses for servers, such as servers 131-135, at which the desired content is mirrored, and returns a list of servers 112 to the client. The client then typically directs its request to the first IP address in the list of servers 112. If no response is received from the server associated with the first IP address, then the client may reissue its request to the second IP address in the list of servers 112 and so on until it finds an available server. In view of this example, it should be appreciated that neither the geographic location of the client nor the geographic location of the server is taken into consideration in determining to which server 131-135 a client should direct a service request.

[0004] **FIG. 2** conceptually illustrates an existing approach for directing requests to the closest server having a shared Anycast address. Anycast addressing is a form of communication that takes place over a network between a client and the “nearest” of a set of servers that can respond to the client’s service request, where “nearest” is determined by network metrics. In this simplified illustration, clients 201-205 are communicatively coupled with servers 231-235 via a network 220. Each of servers 231-235 offers a common service and advertises to router 210 a corresponding shared Anycast address 241-245 of X.X.X.X. When a client, such as one of clients 101-105, issues a request to Anycast address X.X.X.X, such as service request 211, for the service offered by servers 231-235, the router 210 directs the request to the nearest of the servers 231-235 that serves the Anycast address as determined by the most recent network metrics calculated by router 210 or otherwise provided to the router 210. Consequently, a subsequent service request, such as service request 212, even if issued by the same client, may be directed to a different server based on then existing network metrics as observed by router 210. In view of this example, it should be appreciated that Anycast addressing only balances based on current network metrics without regard for the relative load being experienced by servers 231-235. Furthermore, Anycast addressing will not scale beyond the point where the nearest server is incapable of handling all traffic in its area.

[0005] While the load balancing approaches discussed above may be adequate for web traffic and flow-based network communications, they do not address the needs of session-based, latency dependent applications, such as Voice over IP (VoIP), which aspire to reliably minimize latency through the network.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0006] Embodiments of the present invention are 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:

[0007] **FIG. 1** conceptually illustrates an existing Domain Name System (DNS) Round-robin approach for finding an available server offering a desired service.

[0008] **FIG. 2** conceptually illustrates an existing approach for directing requests to the closest server having a shared Anycast address.

[0009] **FIG. 3** conceptually illustrates a high-level geo-locating load balancing architecture according to one embodiment of the present invention.

[0010] **FIG. 4** conceptually illustrates high-level call registration flow according to a redirection embodiment of the present invention.

[0011] **FIG. 5** conceptually illustrates high-level call registration flow according to a redirection by proxying embodiment of the present invention.

[0012] **FIG. 6** conceptually illustrates high-level call registration flow according to a proxy forwarding embodiment of the present invention.

[0013] **FIG. 7** is an example of a computer system with which embodiments of the present invention may be utilized.

[0014] **FIG. 8** is a flow diagram illustrating session establishment processing according to a redirection embodiment of the present invention.

[0015] **FIG. 9** is a flow diagram illustrating session establishment processing according to a redirection by proxying embodiment of the present invention.

[0016] **FIG. 10** is a flow diagram illustrating session establishment processing according to a proxy forwarding embodiment of the present invention.

SUMMARY

[0017] Methods and apparatus are described for geo-locating load balancing. According to one embodiment, the geo-locating load balancing methodology includes a load balancer advertising a virtual Internet Protocol (IP) address shared by one or more other load balancers. The load balancer performs load balancing of requests for services offered by multiple servers corresponding to the load balancer by causing a service request issued by a client to be directed to a particular server.

[0018] In one embodiment, a method is provided for establishing a session for a Voice over IP (VoIP) call. A voice client coupled to a communication network issues a Session Initiation Protocol (SIP) Register message to an Anycast address serviced by multiple proxy servers coupled to the communication network. The SIP Register message is received by the proxy server determined to be closest to the voice client based on metrics associated with the communication network. The closest proxy server then causes the SIP Register message to be directed to a particular registrar server of multiple registrar servers associated with the proxy server based on a load balancing routine.

[0019] According to one embodiment, a novel communication network architecture is provided, including multiple servers, multiple load balancers, and multiple geographically dispersed communication devices. The servers provide services to the communication devices within the communication network. The load balancers each service a shared virtual Internet Protocol (IP) address common to all of the load balancers and perform load balancing of service requests on behalf of two or more of the servers that are located geographically proximate to the load balancer. The communication devices are communicatively coupled with the load balancers and are configured to issue service requests intended for any of the servers to the shared virtual IP address. In this manner, upon issuing a service request, a communication device is directed to a particular server selected by a load balancing routine that is associated with a load balancer that is closest to the communication device.

[0020] Other features of embodiments of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

DETAILED DESCRIPTION

[0021] Methods and apparatus are described for geo-locating load balancing. Broadly stated, embodiments of the present invention seek to provide a mechanism for directing geographically dispersed clients to the "closest" feature servers without previously determining their locations and then balancing the load across the closest feature servers. According to one embodiment, a method is provided for establishing a session for a Voice over IP (VoIP) call. A voice client coupled to a communication network issues a Session Initiation Protocol (SIP) Register message to an Anycast address serviced by multiple proxy servers coupled to the communication network. The SIP Register message is received by the proxy server determined to be closest to the voice client based on metrics associated with the communication network. The closest proxy server then causes the SIP Register message to be directed to a particular registrar server of multiple registrar servers associated with the proxy server based on a load balancing routine.

[0022] Depending upon the implementation the load balancer may operate in accordance with one or more service request processing methodologies referred to herein as "redirection," "redirection by proxying," and "proxy forwarding." When performing redirection and redirection by proxying, the load balancer causes the service request to be directed to the particular server and is thereafter excluded from subsequent messaging flow exchanged between the particular server and the client relating to the session. When

performing proxy forwarding, the load balancer remains within subsequent messaging flow exchanged between the particular server and the client relating to the session.

[0023] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of embodiments of the present invention. It will be apparent, however, to one skilled in the art that embodiments of the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

[0024] Embodiments of the present invention include various steps, which will be described below. The steps may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware, software, and/or firmware.

[0025] Embodiments of the present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, compact disc read-only memories (CD-ROMs), and magneto-optical disks, ROMs, random access memories (RAMs), erasable programmable read-only memories (EPROMs), electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, flash memory, or other type of media / machine-readable medium suitable for storing electronic instructions. Moreover, embodiments of the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

[0026] While, for convenience, embodiments of the present invention may be described with reference to Internet Protocol (IP) originating voice products/services and other VoIP applications, the present invention is equally applicable to various other session-based, latency dependent applications and/or applications that require real-time performance, such as online gaming, instant messaging, applications based on human interactions (e.g., collaborative software, online/Web collaboration, voice conferencing, and video conferencing), and real-time data communication and/or exchange, such as market data applications, financial transactions, and the like.

Terminology

[0027] Brief definitions of terms used throughout this application are given below.

[0028] The terms “closeness,” “nearness,” “closest,” “nearest” and the like are used herein in a logical sense and are not necessarily limited to physical proximity. According to one embodiment, one or more or a combination of various network metrics, such as link congestion, cost metrics assigned to given routers, etc., are relied upon to determine the closeness or nearness (e.g., the logical proximity) of two devices communicatively coupled to a packet-based network. In some cases, the logical proximity between devices is determined by or based upon link-state information and/or routing protocols, such as Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS) or the like. In one embodiment, the closest or nearest of a plurality of load balancers to a particular client is the load balancer with which the client is capable of communicating and experiencing the least latency and/or traversing the fewest hops. Depending on the context, this notion of logical proximity may have different meanings. For example, in the context of a client connected directly to the target service network in which the novel load balancers reside, the client may be directed to the geographically closest operating load balancer based on routing metrics within the target service network. In the context of a client connected indirectly to the target service network in which the novel load balancers reside (via a border network, for example), the client may be directed to the load balancer geographically closest to that Internet Service Provider’s (ISP’s) preferred point of interconnection with the target service network.

[0029] The phrase “communication device” or the term “client” generally refers to a device whereby communications or other information are directly or indirectly introduced to or received from a communication network. Thus, as just some examples, communication devices may include, but are not limited to, IP phones, H.323 phones, Session Initiation Protocol (SIP) phones, VoIP phones, Terminal Adapters (TAs), Analog Terminal Adapters (ATAs), Personal Digital Assistants (PDAs), cellular or mobile phones, Personal Computers (PCs), Digital Subscriber Line (DSL) modems, dial up modems, cable modems and the like.

[0030] The phrase “communication network” or term “network” generally refers to a group of interconnected devices capable of exchanging information. A communication network may be as few as several personal computers on a Local Area Network (LAN) or as large as the Internet, a worldwide network of computers. As used herein “communication network” is intended to encompass any network capable of transmitting information from one

entity to another. In one particular case, a communication network is a Voice over Internet Protocol (VoIP) network. In some cases, a communication network may be comprised of multiple networks, even multiple heterogeneous networks, such as one or more border networks, voice networks, broadband networks, service provider networks, Internet Service Provider (ISP) networks, and/or Public Switched Telephone Networks (PSTNs), interconnected via gateways operable to facilitate communications between and among the various networks.

[0031] The terms “connected” or “coupled” and related terms are used in an operational sense and are not necessarily limited to a direct or physical connection or coupling.

[0032] The phrase “feature server” generally refers to a server that is operable to provide one or more services supported by a communications network, such as a voice network. For example, a feature server may provide telecommunications services, such as caller identification, call forwarding, voice mail, and/or the like. In one embodiment, a feature server comprises a Class-5 soft switch. In other embodiments, a feature server may represent, a registrar server.

[0033] The phrase “registrar server” generally refers to a particular type of feature server that performs registration and/or call routing. According to one embodiment, a registrar server processes and maintain SIP Registrations for each of the call parties to enable it to route calls when a SIP INVITE is received. For example, when one user wishes to call another, both clients may perform a SIP Registration with the registrar server; this provides enough information for the registrar server to route the call when the INVITE is later sent from one user to the other.

[0034] The phrases “in one embodiment,” “according to one embodiment,” and the like generally mean the particular feature, structure, or characteristic following the phrase is included in at least one embodiment of the present invention, and may be included in more than one embodiment of the present invention. Importantly, such phrases do not necessarily refer to the same embodiment.

[0035] If the specification states a component or feature “may”, “can”, “could”, or “might” be included or have a characteristic, that particular component or feature is not required to be included or have the characteristic.

[0036] The phrase “load balancer” generally refers to a logical or physical device that performs load balancing, such as hardware load balancing solutions performed by computer

systems, switches and/or routers, software load balancing solutions, load balancing servers and the like. According to one embodiment load balancer software is run on an Edge Proxy Server, such as a Netra 240 Server available from Sun Microsystems. According to an embodiment, load balancing may be performed by one or more dedicated hardware-based flow-based load balancers which selectively send traffic from clients to particular servers based on the amount of load that the load balancer is currently sending to the servers for which it fronts. According to one embodiment, a load balancer may be capable of forwarding to additional devices in front of other servers or server farms if the servers for which it fronts are too heavily loaded to answer requests. Depending upon the particular implementation and the type of services to be supported, a load balancer may be stateful, stateless, or semi-stateful.

[0037] The phrase “load balancing” generally refers to a method of taking multiple requests or processes and distributing each of them across multiple computers or network devices. According to one embodiment the distribution among the multiple computers or network devices is based on how busy the computer or network device is or based on historical information regarding how previous requests or processes were distributed among the devices.

[0038] The term “responsive” includes completely or partially responsive.

[0039] **FIG. 3** conceptually illustrates a high-level geo-locating load balancing architecture according to one embodiment of the present invention. In an attempt to address various performance issues and other deficiencies of existing load balancing methodologies, the load balancing architecture described herein associates a virtual IP address, such as a shared Anycast address, with multiple load balancers residing in front of corresponding server farms or clusters. The load balancers each advertise the virtual IP address to the routing protocol of the communication network, such as the OSPF routing protocol. This allows service requests to be directed to the closest load balancer based on existing network metrics. For example, in the context of a communications network implementing the OSPF routing protocol, service requests directed to a shared Anycast address will reach the closest load balancer based on OSPF rules. The closest load balancer may then balance the load across the feature servers within the corresponding server farm in accordance with customary load balancing techniques. For example, the feature servers within a server farm or cluster may provide periodic status updates to the load balancer associated with that server farm or cluster to allow the load balancer to distribute service requests appropriately across the feature

servers. According to one embodiment, in a SIP environment, a load balancer may use a SIP OPTIONS message to test liveliness of feature servers. The feature servers may provide load feedback to the load balancer via a SIP OPTIONS response message.

[0040] In this simplified illustration, each of a plurality of distinct geographic regions, such as region 330, region 340 and region 350, includes at least one load balancer and at least one server. Clients 301-305 are communicatively coupled through network 320 with server farms 331, 341 and 351 via load balancers 335, 345 and 365, respectively.

[0041] According to this example, clients 301-305 are provisioned with a Uniform Resource Identifier (URI), a Uniform Resource Locator (URL) or domain name, such as relay.level3.com, which will resolve to an address used as a virtual IP address 336 on all of the load balancers 335, 345 and 365. In alternative embodiments, clients 301-305 may be provisioned with one or more virtual IP addresses. When a client needs access to or would like to register with a particular service offered by one of server farms 331, 341 and 351, it issues a DNS name 311 request, such as a Domain Name Service lookup, containing the provisioned domain name. A DNS server 310 translates the domain name within the DNS name 311 request into the corresponding virtual IP address 336 and returns the virtual IP address to the requesting client via virtual IP address 312 message.

[0042] Continuing with the present example, the client may then direct a service request 321 to the virtual IP address 336 returned by the DNS server 310. A router 320 receiving the service request 321 then directs the service request 321 to the closest load balancer (e.g., load balancer 365) based on its internal routing tables which have been populated in accordance with routing protocol rules, such as OSPF rules. The load balancer 365, then redirects the request to an appropriate server (not shown) within the server farm 351. This redirection may be distributed among the servers within the server farm 351 based on past redirection, in a round-robin fashion, or based on load feedback communicated to the load balancer by the servers of the server farm 341. Various other current or future load balancing methodologies may be employed, such as least connections, weighted round robin, weighted least connections, address space hashing and/or URL hashing. According to one embodiment, once a service request, such as a SIP REGISTER message, has reached the closest load balancer, the load balancer will reply with a SIP 302 Temporarily Moved message containing the next entry in a rotating list of addresses of clusters of Real-Time Transport Protocol (RTP) relay devices. Advantageously, in this manner, geographically

dispersed clients may be directed to the closest feature servers while the load is balanced appropriately among such feature servers.

[0043] In the simplified examples discussed herein, it should be appreciated that the communication networks may be comprised of multiple heterogeneous networks and clients, load balancers, feature servers, Network Address Translation (NAT) traversal managers (NTMs) and other network devices may be associated with the same network or different networks within a particular communication network. Furthermore, while a single router, such as router 320, is depicted in various examples described herein, it should be appreciated that messaging exchanged between clients and load balancers, feature servers and/or NTMs need not be via a particular router or set of routers. Rather, a router is depicted in the logical sense merely to convey the fact that routing protocols and/or network metrics are involved in the process of directing client service requests to the nearest load balancer. In a real world scenario, messaging exchanged among clients, load balancers, features servers, NTMs and the like may traverse many different routers.

[0044] **FIG. 4** conceptually illustrates high-level call registration flow according to a redirection embodiment of the present invention. According to one embodiment, the present redirection methodology is used to initially register a voice client to a stateless load balancer, which redirects the client directly to a feature server or indirectly to a feature server via an intermediate NTM. In such an embodiment, after the initial registration flow, the load balancer is no longer part of the rest of the SIP messaging flow between the client and the NTM and/or the feature server.

[0045] According to the present example, prior to initiating a VoIP call or periodically, a client 401, such as a VoIP phone provisioned with a domain name, is configured to issue a DNS name 411 request to DNS server 410 to obtain an IP address to which a service request 421a, such as a registration request, is to be directed.

[0046] The DNS server 410 maintains a mapping of domain names to IP addresses and returns a virtual IP address 412 message containing a virtual IP address 431 of load balancer 430 which is shared by one or more other geographically distributed load balancers (not shown).

[0047] Continuing with the present example, after receiving the shared virtual IP address 431 from the DNS server 410, the client 401 may then direct a service request 421a to the shared virtual IP address 431. As above, a router 420 receiving the service request 421a

directs the service request 421b to the closest load balancer (e.g., load balancer 430) based on routing protocol rules.

[0048] Load balancer 430 is operable to redirect a communication device, such as client 401, to an appropriate feature server, such as server 440 or server 450. In some cases, the load balancer 430 may be a proxy server configured to balance the load among a plurality of feature servers based on a load balancing routine.

[0049] According to the present example, the load balancer 430 redirects client 401 to server 440 by transmitting a redirect message 422a which is delivered to client 401 via router 420 as redirect message 422b. According to one embodiment, the redirect message 422a, 422b contains a unicast address 441 associated with server 440 thereby removing load balancer 430 from subsequent messaging flow by communicating to client 401 that subsequent requests should be addressed directly to server 440 using unicast address 441. According to one embodiment, the redirect message 422a, 422b comprises a SIP Moved Temporarily Message with one or more redirect URIs.

[0050] Continuing with the present example, after receiving the redirect message 422b, the client 401 is configured to direct subsequent messaging flow and/or requests, such as request 423a, to unicast address 441 which are delivered to server 440 via router 420.

[0051] **FIG. 5** conceptually illustrates high-level call registration flow according to a redirection by proxying embodiment of the present invention. According to one embodiment, the present redirection by proxying methodology is used to register a client to a stateful or semi-stateful load balancer which remains in the SIP messaging flow even after the initial registration flow. After the load balancer redirects the initial registration request to an appropriate feature server or NTM-feature server pair, it may maintain call context to redirect subsequent SIP messaging flow and/or requests from the client to the same feature server or NTM-feature server pair. Alternatively, in a non-session-based application environment, the load balancer may be stateless as there is no need to ensure requests are redirected to the same feature server or NTM-feature server pair.

[0052] According to the present example, prior to initiating a VoIP call or periodically, a client 501, such as a VoIP phone provisioned with a domain name, is configured to issue a DNS name 511 request to DNS server 510 to obtain an IP address to which a service request 521a, such as a registration request, is to be directed.

[0053] The DNS server 510 maintains a mapping of domain names to IP addresses and is configured to translate the domain name within the DNS name 511 request into the

corresponding IP address and return a virtual IP address 512 message containing a virtual IP address 531 of load balancer 530 which is shared by one or more other geographically distributed load balancers (not shown).

[0054] Continuing with the present example, after receiving the shared virtual IP address 531 from the DNS server 510, the client 501 may then direct a service request 521a to the shared virtual IP address 531. As above, a router 520 receiving the service request 521a directs the service request 521b to the closest load balancer (e.g., load balancer 530) based on routing protocol rules.

[0055] Load balancer 530 is operable to perform redirection by proxying to redirect a communication device, such as client 501, to an appropriate feature server, such as server 540 or server 550. In some cases, the load balancer 530 may be a proxy server configured to balance the load among a plurality of feature servers based on a load balancing routine.

[0056] According to the present example, the load balancer 530 redirects client 501 to server 540 by forwarding request 521c to server 540. Servers 540 and 550 are configured to direct their responses to the client 501 and identify themselves as the source of the responses, such as response 522a which is delivered to client 501 via router 520 as response 522b. According to one embodiment, the response 522a, 522b, contains a unicast address 541 associated with server 540 thereby removing load balancer 530 from subsequent messaging flow by communicating to client 501 that subsequent requests should be addressed directly to server 540 using unicast address 541. Alternatively, the response 522a, 522b may include a SIP URI associated with server 540.

[0057] Continuing with the present example, after receiving the response 522b, the client 501 is configured to direct subsequent messaging flow and/or requests, such as request 523a, to unicast address 541 which are delivered to server 540 via router 520.

[0058] **FIG. 6** conceptually illustrates high-level call registration flow according to a proxy forwarding embodiment of the present invention. According to one embodiment, the present proxy forwarding methodology is used to register a client to a stateful or semi-stateful load balancer which remains in the SIP messaging flow even after the initial registration flow. After the load balancer forwards the initial registration request to an appropriate feature server or NTM-feature server pair, it may maintain call context to redirect subsequent SIP messaging flow and/or requests from the client to the same feature server or NTM-feature server pair. Alternatively, in a non-session-based application environment, the proxy

forwarding load balancer may be stateless as there is no need to ensure requests are forwarded to the same feature server or NTM-feature server pair.

[0059] According to the present example, prior to initiating a VoIP call or periodically, a client 601, such as a VoIP phone provisioned with a domain name, is configured to issue a DNS name 611 request to DNS server 610 to obtain an IP address to which a service request 621a, such as a registration request, is to be directed.

[0060] The DNS server 610 maintains a mapping of domain names to IP addresses and is configured to translate the domain name within the DNS name 611 request into the corresponding IP address and return a virtual IP address 612 message containing a virtual IP address 631 of load balancer 630 which is shared by one or more other geographically distributed load balancers (not shown).

[0061] As above, after receiving the shared virtual IP address 631 from the DNS server 610, the client 601 may then direct a service request 621a to the shared virtual IP address 631 which is directed to the closest load balancer (e.g., load balancer 630) as request 621b based on routing protocol rules.

[0062] Load balancer 630 is operable to perform proxying by forwarding to deliver requests from communication devices, such as client 601, to an appropriate feature server, such as server 640 or server 650. In some cases, the load balancer 630 may be a proxy server configured to balance the load among a plurality of feature servers based on a load balancing routine.

[0063] According to the present example, the load balancer 630 forwards request 621b as request 621c to server 640. Servers 640 and 650 are configured to direct their responses to the load balancer 630 which in turn is operable to forward such responses to the requesting client and identify itself as the source of the responses, such as response 622a, 622b which is delivered to client 601 via router 620 as response 622c. According to one embodiment, the response 622b, 622c, contains a unicast address 632 associated with the load balancer 630 thereby maintaining load balancer 630 within subsequent messaging flow by communicating to client 601 that subsequent requests should be addressed to load balancer 630 using unicast address 632. Alternatively, the response 622b, 622c may include a SIP URI associated with load balancer 630.

[0064] Continuing with the present example, after receiving the response 622c, the client 601 is configured to direct subsequent messaging flow and/or requests to unicast address 632 which are delivered to load balancer 630 via router 620.

[0065] **Figure 7** is an example of a computer system 700 with which embodiments of the present invention may be utilized. Computer system 700 represents an exemplary load balancer or proxy server which may implement one or more of the redirection or forwarding mechanisms described herein (i.e., redirection, redirection by proxying and proxy forwarding). In this simplified example, the computer system 700 comprises a bus 701 or other communication means for communicating data and control information, and one or more processors 702, such as Intel® Itanium® or Itanium 2 processors or Sun® UltraSPARC-III® processors, coupled with bus 701.

[0066] Computer system 700 further comprises a random access memory (RAM) or other dynamic storage device (referred to as main memory 704), coupled to bus 701 for storing information and instructions to be executed by processor(s) 702. Main memory 704 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor(s) 702.

[0067] Computer system 700 also comprises a read only memory (ROM) 706 and/or other static storage device coupled to bus 701 for storing static information and instructions for processor(s) 702.

[0068] A mass storage device 707, such as a magnetic disk or optical disc and its corresponding drive, may also be coupled to bus 701 for storing instructions and information, such as configuration files, a key store and registration database, etc.

[0069] One or more communication ports 703 may also be coupled to bus 701 for supporting network connections and communication of information to/from the computer system 700 by way of a communication network, such as a Local Area Network (LAN), Wide Area Network (WAN), the Internet, or PSTNs, for example. The communication ports 703 may include various combinations of well-known interfaces, such as one or more modems to provide dial up capability, one or more 10/100 Ethernet ports, one or more Gigabit Ethernet ports (fiber and/or copper), or other well-known network interfaces commonly used in internetwork environments. In any event, in this manner, the computer system 700 may be coupled to a number of other network devices, communication devices, clients, NTMs, and/or servers via a conventional communication network infrastructure.

[0070] Optionally, operator and administrative interfaces (not shown), such as a display, keyboard, and a cursor control device, may also be coupled to bus 701 to support direct operator interaction with computer system 700. Other operator and administrative

interfaces can be provided through network connections connected through communication ports 703.

[0071] Finally, removable storage media (not shown), such as one or more external or removable hard drives, tapes, floppy disks, magneto-optical discs, compact disk-read-only memories (CD-ROMs), compact disk writable memories (CD-R, CD-RW), digital versatile discs or digital video discs (DVDs) (e.g., DVD-ROMs and DVD+RW), Zip disks, or USB memory devices, e.g., thumb drives or flash cards, may be coupled to bus 701 via corresponding drives, ports or slots.

[0072] **FIG. 8** is a flow diagram illustrating session establishment processing according to a redirection embodiment of the present invention. In the present example, the client is assumed to have knowledge of a shared virtual IP address associated with multiple geographically distributed load balancers. As discussed above with reference to FIGS. 4-6, the client may obtain the shared virtual IP address in various ways. For example, the client may obtain the shared virtual IP address during a provisioning or configuration process and/or may receive the shared virtual IP address responsive to a DNS lookup based on a provisioned domain name associated with the shared virtual IP address.

[0073] At block 810, a client (e.g., voice client 401) issues a service request (e.g., service request 421a) to the shared virtual IP address. According to one embodiment, the service request is part of an initial registration flow to establish a session between the client and a feature server or between the client and another communication device.

[0074] At block 820, the service request is routed to the closest load balancer advertising the shared virtual IP address. According to one embodiment, multiple geographically distributed load balancers advertise the shared virtual IP address to the OSPF routing protocol thereby defining closeness in terms of the OSPF rules. In alternative embodiments, closeness may have different meanings. For example, in the context of a client connected directly to the target service network in which the novel load balancers reside, the client may be directed to the geographically closest operating load balancer based on routing metrics within the target service network. In the context of a client connected indirectly to the target service network in which the novel load balancers reside (via a border network, for example), the client may be directed to the load balancer geographically closest to that Internet Service Provider's (ISP's) preferred point of interconnection with the target service network.

[0075] At block 830, the load balancer (e.g., load balancer 430) redirects the client to a server or an NTM-server pair in accordance with its load balancing methodology. In the context of a client registering to participate in SIP-based voice conversations, for example, the redirection may include a SIP Moved Temporarily Message with one or more redirect URIs or an indication of a unicast address (e.g., unicast address 441) associated with the appropriate server or NTM.

[0076] According to one embodiment, load balancers and feature servers (e.g., registrar servers) or load balancers and NTM-feature server pairs are geographically distributed in physical proximity to each other to ensure clients are redirected to a feature server or NTM-feature server pair that is relatively close to the client making the request.

[0077] At block 840, a session is established between the client and the server (e.g., server 440) or NTM to which it was redirected. According to one embodiment, in response to a redirect message (e.g., redirect 422a, 422b) the client is configured to direct subsequent messaging related to the session, such as SIP messaging flow, to the server or NTM-server pair identified by the redirect message.

[0078] Advantageously, in this manner, geographically dispersed clients, such as voice clients, may be directed to the nearest feature servers while also balancing the load across the feature servers. Another desirable effect of the various geo-locating load balancing models described herein is the load balancer can be established as the only place where NTM traffic distribution occurs. Hence, provisioning may be centralized rather than potentially being spread out over millions of communication devices. The above described redirection embodiment also has the benefit of being able to be implemented by a stateless load balancer or proxy server that can be streamlined for performing efficient redirect functionality. Finally, because the load balancers need only participate in the initial register flow and can be excluded from the rest of the messaging flow, the load balancers can scale independently from the feature servers or the NTM-feature server pairs.

[0079] **FIG. 9** is a flow diagram illustrating session establishment processing according to a redirection by proxying embodiment of the present invention. As above, in the present example, the client is assumed to have knowledge of a shared virtual IP address associated with multiple geographically distributed load balancers. The particular mechanism for determining or becoming aware of the shared virtual IP address is not of particular importance. Rather, it is the fact that multiple load balancers share and advertise a common virtual IP address that ultimately enables the general notion of geo-locating load balancing.

[0080] At block 910, a client (e.g., voice client 501) issues a service request (e.g., service request 521a) to the shared virtual IP address. According to one embodiment, the service request is part of an initial registration flow to establish a session between the client and a feature server or between the client and another communication device.

[0081] At block 920, the service request is routed to the closest load balancer advertising the shared virtual IP address. According to one embodiment, multiple geographically distributed load balancers advertise the shared virtual IP address to the OSPF routing protocol thereby defining closeness in terms of the OSPF rules. In alternative embodiments, one or more or a combination of various network metrics, such as link congestion, cost metrics assigned to given routers, etc., may be relied upon to determine the closeness or nearness (e.g., the logical proximity) of two devices communicatively coupled to a communication network. In some cases, the logical proximity between devices is determined by or based upon link-state information and/or routing protocol rules. In one embodiment, the closest or nearest of a plurality of load balancers to a particular client is the load balancer with which the client is capable of communicating and experiencing the least latency and/or traversing the fewest hops.

[0082] At block 930, the load balancer (e.g., load balancer 530) redirects the client to a server or an NTM-server pair in accordance with its load balancing methodology. In the context of a client registering to participate in SIP-based voice conversations, for example, the redirection may be indicated by a response to a SIP Register request (e.g., response 522a) including one or more URIs associated with the appropriate server (e.g., server 540) or an indication of a unicast address (e.g., unicast address 541) associated with the appropriate server or NTM.

[0083] According to one embodiment, load balancers and feature servers (e.g., registrar servers) or load balancers and NTM-feature server pairs are geographically distributed in physical proximity to each other to ensure clients are redirected to a feature server or NTM-feature server pair that is relatively close to the client making the request.

[0084] At block 940, a session is established between the client and the server (e.g., server 540) or NTM to which the client was redirected. According to one embodiment, in response to receiving the redirection indication (e.g., response 522b) the client is configured to direct subsequent messaging related to the session, such as SIP messaging flow, to the server or NTM-server pair identified by the message containing or otherwise conveying the redirection indication.

[0085] **FIG. 10** is a flow diagram illustrating session establishment processing according to a proxy forwarding embodiment of the present invention. As above, in the present example, the client is assumed to have knowledge of a shared virtual IP address associated with multiple geographically distributed load balancers.

[0086] At block 1010, a client (e.g., voice client 601) issues a service request (e.g., service request 621a) to the shared virtual IP address. According to one embodiment, the service request is part of an initial registration flow to establish a session between the client and a feature server or between the client and another communication device.

[0087] At block 1020, the service request is routed to the closest load balancer advertising the shared virtual IP address as described earlier.

[0088] At block 1030, the load balancer (e.g., load balancer 630) forwards the service request (e.g., request 621c) to a server or an NTM-server pair in accordance with its load balancing methodology.

[0089] At block 1040, the server or NTM-server pair to which the service request was forwarded responds (e.g., response 622a) to the load balancer.

[0090] At block 1050, the load balancer forwards the response (e.g., response 622b) to the client and directs the client to transmit subsequent messaging flow relating to this session to the load balancer. According to one embodiment, the load balancer directs the client to transmit subsequent messaging flow relating to this session to the load balancer by identifying the load balancer's unicast address as the source of the forwarded response (e.g., response 622b). Alternatively, the forwarded response may include a SIP URI associated with load balancer 630.

[0091] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A method comprising:
a load balancer advertising a virtual Internet Protocol (IP) address shared by one or
5 more other load balancers, the load balancer performing load balancing of requests for
services offered by a plurality of servers corresponding to the load balancer; and
responsive to receiving a service request from a client, the load balancer causing
the service request to be directed to a particular server of the plurality of servers.
- 10 2. The method of claim 1, wherein the virtual IP address comprises an
Anycast address.
3. The method of claim 1, wherein the load balancing is based upon the
relative loads being experienced by the plurality of servers as reported to the load balancer
by the plurality of servers.
- 15 4. The method of claim 1, wherein the load balancer causing the service
request to be directed to a particular server comprises the load balancer redirecting the
client to the particular server and the load balancer is excluded from subsequent
messaging flow exchanged between the particular server and the client.
- 20 5. The method of claim 1, wherein the load balancer causing the service
request to be directed to a particular server comprises the load balancer performing
redirection by proxying and the load balancer is excluded from subsequent messaging
flow exchanged between the particular server and the client.
- 25 6. The method of claim 1, wherein the load balancer causing the service
request to be directed to a particular server comprises the load balancer performing proxy
forwarding and the load balancer remains within subsequent messaging flow exchanged
between the particular server and the client.
7. The method of claim 1, wherein the load balancer and the one or more
other load balancers are physically located within a common geographic region.
8. The method of claim 1, wherein the client is provisioned with a domain
name associated with the virtual IP address, and wherein the method further comprises:
30 prior to the load balancer receiving the service request, the client issuing a
translation request to a Domain Name System (DNS) server for a translation of the
provisioned domain name; and

responsive to the translation request the DNS server returning the virtual IP address the client.

9. The method of claim 1, wherein the service request is associated with a Voice over IP (VoIP) call.

5 10. The method of claim 9, wherein the service request comprises a Session Initiation Protocol (SIP) Register message.

11. The method of claim 10, wherein the particular server comprises a feature server.

10 12. The method of claim 10, wherein the particular server comprises a registrar server.

13. The method of claim 1, wherein the method is performed as a result of one or more processors executing instructions stored on a computer-readable medium.

14. A method of establishing a session for a Voice over IP (VoIP) call comprising:

15 a voice client coupled to a communication network issuing a Session Initiation Protocol (SIP) Register message to an Anycast address serviced by a plurality of proxy servers coupled to the communication network;

20 the SIP Register message being received by a proxy server of the plurality of proxy servers determined to be closest to the voice client based on metrics associated with the communication network; and

the proxy server causing the SIP Register message to be directed to a particular registrar server of a plurality of registrar servers associated with the proxy server based on a load balancing routine.

25 15. The method of claim 14, wherein the method is performed as a result of one or more processors executing instructions stored on a computer-readable medium.

16. A method comprising:

30 responsive to receiving a service request from a voice client of a plurality of geographically dispersed voice clients coupled to a communication network, causing the service request to be transmitted to a load balancer associated with a closest set of feature servers coupled to the communication network; and

causing the service request message to be directed to a particular feature server of the closest set of feature servers based on a load balancing routine.

17. The method of claim 16, wherein the method is performed as a result of one or more processors executing instructions stored on a computer-readable medium.

18. A communication network comprising:

5 a plurality of servers providing services to communication devices within the communication network;

a plurality of load balancers each servicing a shared virtual Internet Protocol (IP) address common to the plurality of load balancers, each of the plurality of load balancers performing load balancing of service requests on behalf of two or more of the plurality of servers located geographically proximate to the load balancer; and

10 a plurality of geographically dispersed communication devices communicatively coupled with the plurality of load balancers, each of the plurality of geographically dispersed communication devices configured to issue service requests intended for any of the plurality of servers to the shared virtual IP address, whereby upon issuing a service request, a communication device of the plurality of geographically dispersed
15 communication devices is directed to a particular server of the plurality of servers that is associated with a load balancer of the plurality of load balancers that is closest to the communication device, the particular server selected by a load balancing routine executing on the load balancer.

19. The communication network of claim 18, wherein each of the plurality of
20 servers comprise feature servers.

20. The communication network of claim 18, wherein each of the plurality of load balancers comprise proxy servers.

21. The communication network of claim 18, wherein one or more of the plurality of geographically dispersed communication devices comprise Voice over IP
25 (VoIP) phones.





FIG. 1 (Prior Art)



200

FIG. 2 (Prior Art)

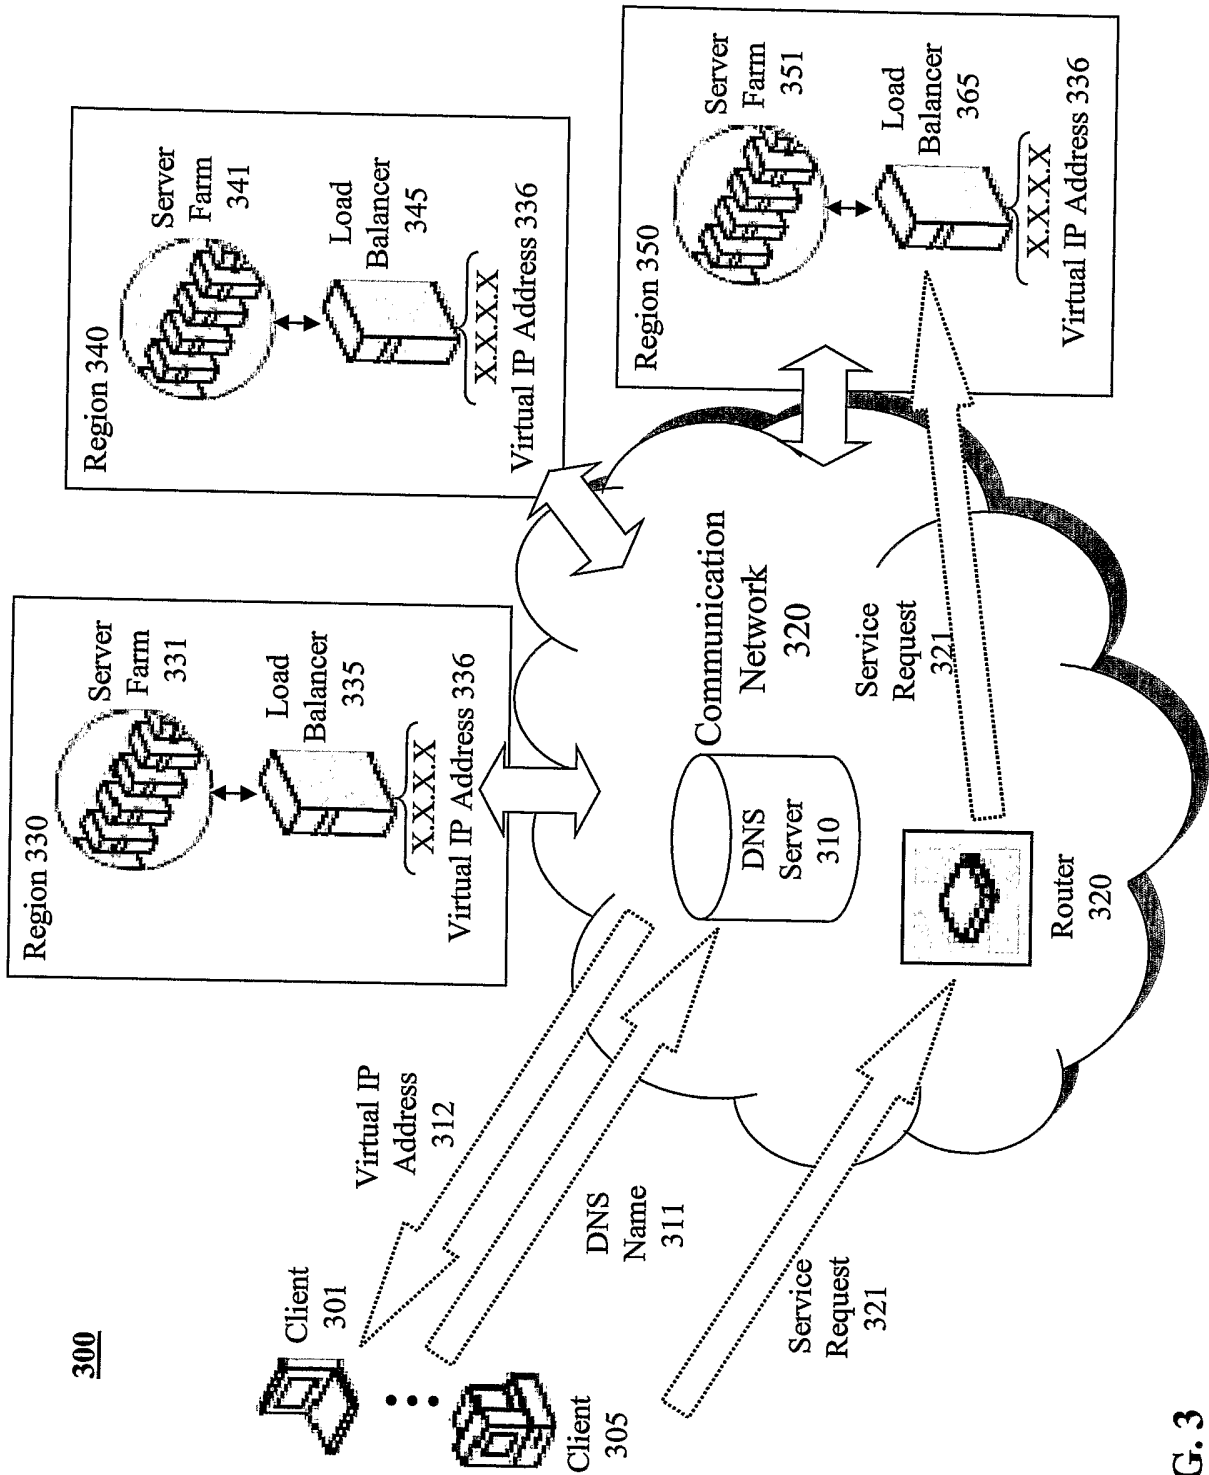


FIG. 3




FIG. 4




FIG. 5

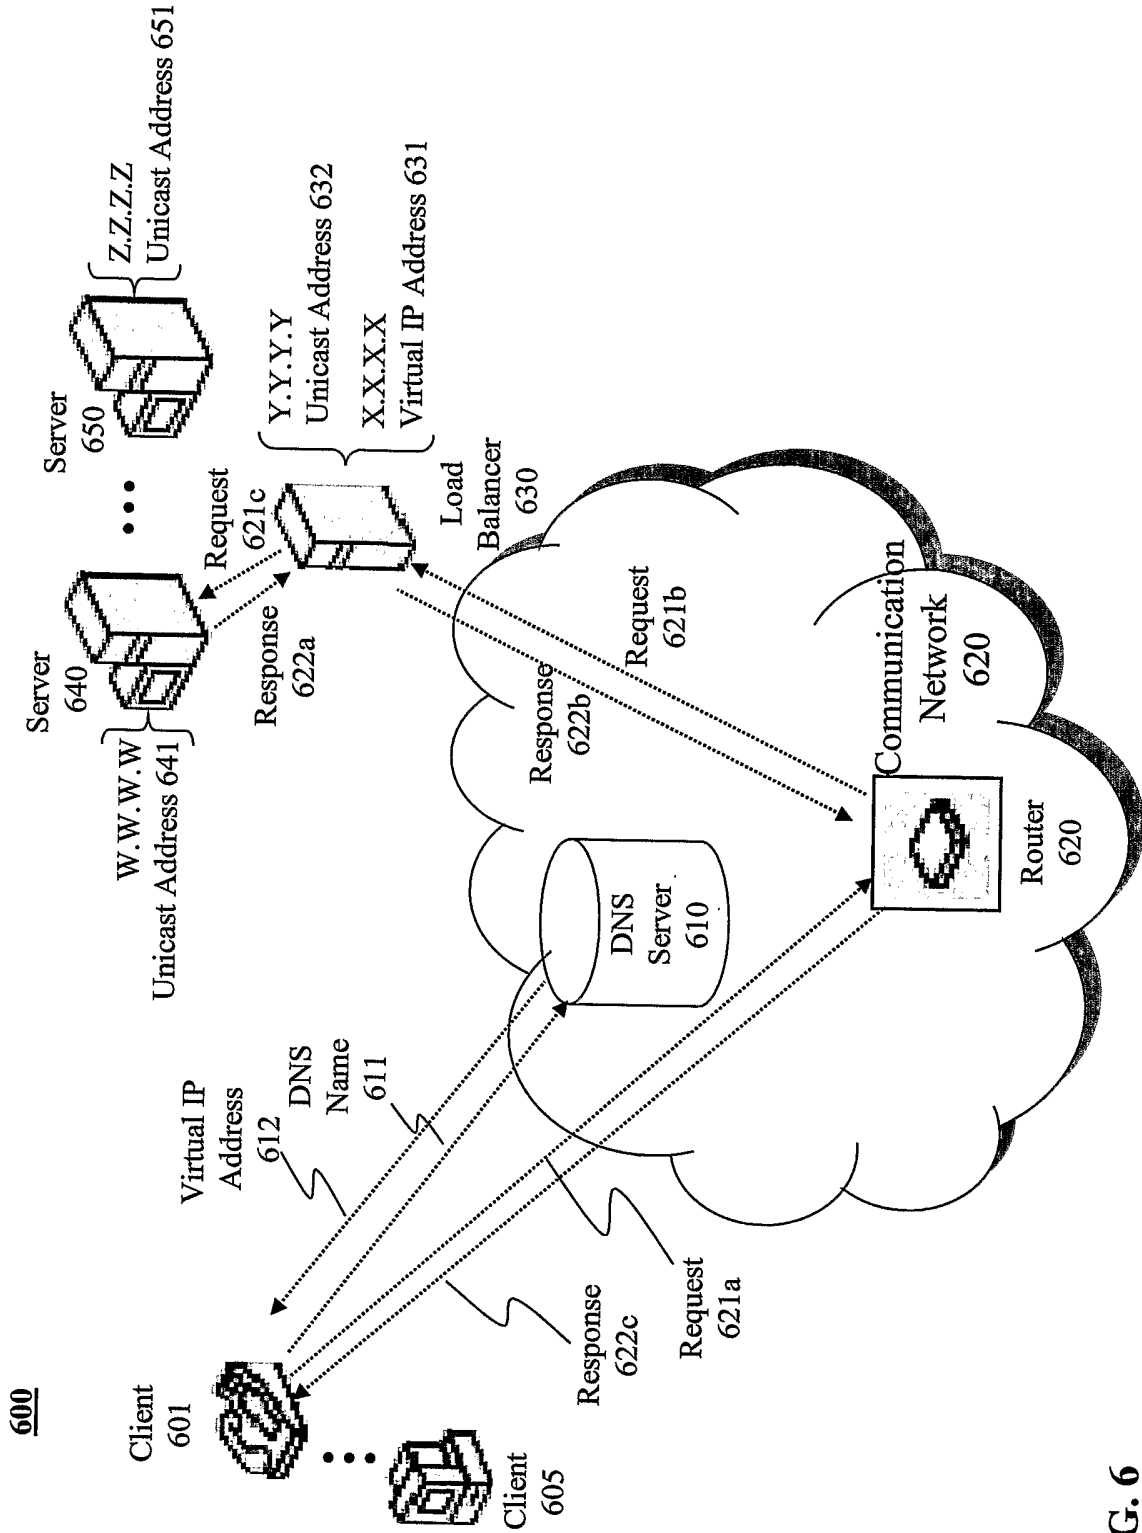



FIG. 6

7/10



700

FIG. 7

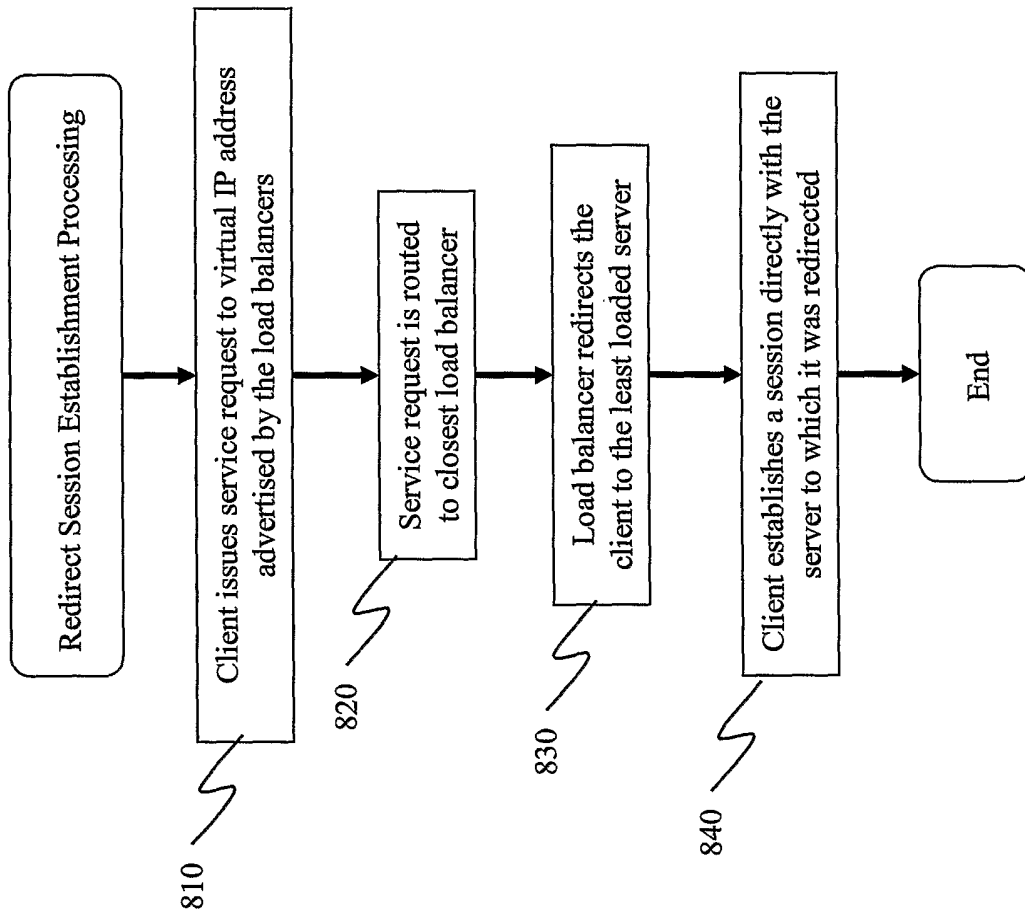


FIG. 8

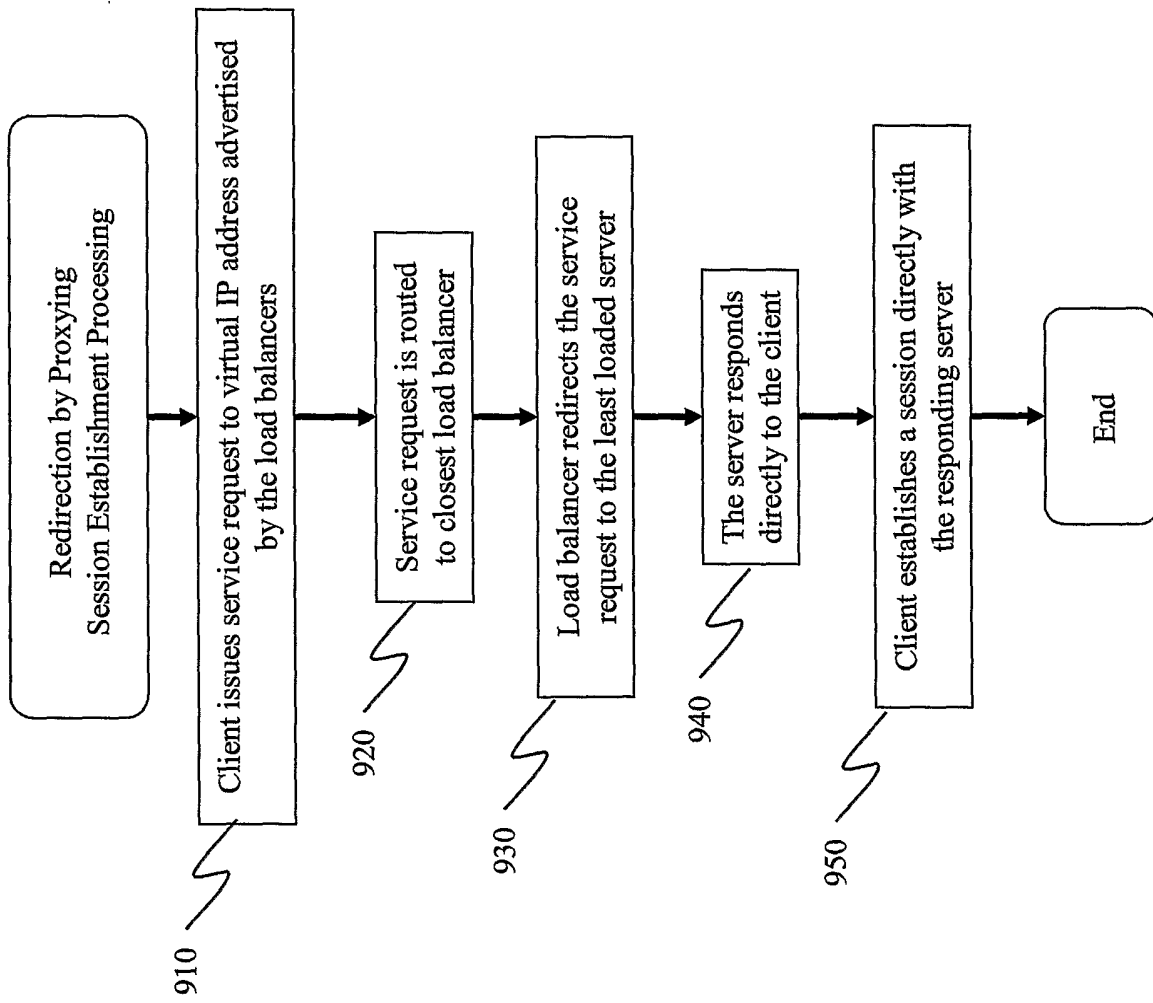


FIG. 9

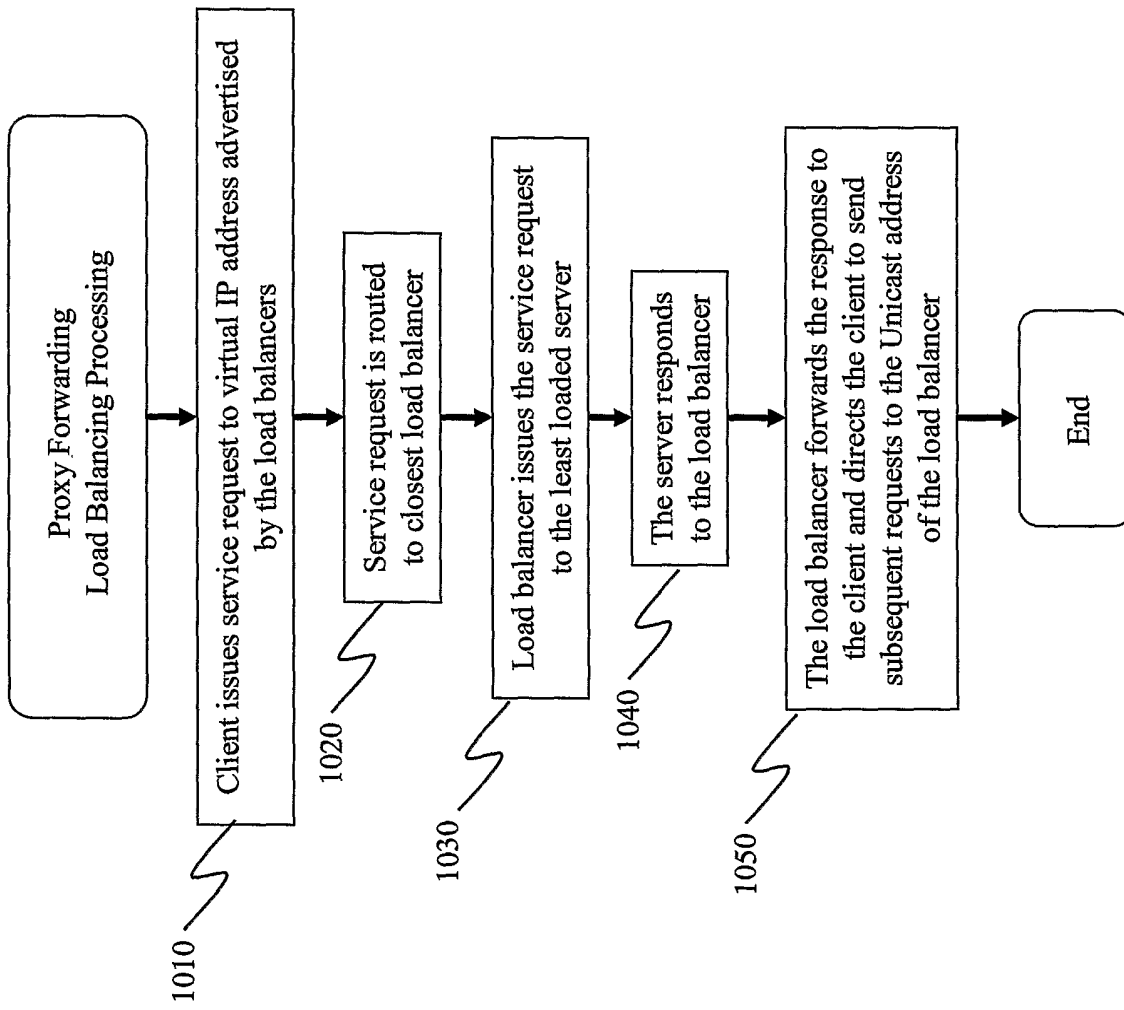


FIG. 10