



- (51) **International Patent Classification:**
H04L 29/08 (2006.01) **G06F 15/173** (2006.01)
G06F 9/50 (2006.01)
- (21) **International Application Number:**
PCT/US2016/034493
- (22) **International Filing Date:**
26 May 2016 (26.05.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/179,969 26 May 2015 (26.05.2015) US
- (71) **Applicant: IDEVICES, LLC** [US/US]; 136 Simsbury Road, Building 12, Avon, CT 06001 (US).
- (72) **Inventor: DJAKOVIC, Vladan;** 601 Van Ness Avenue, E-841, San Francisco, CA 94102 (US).
- (74) **Agents: REINER, Kevin, L. et al.;** McCarter & English, LLP, CityPlace I, 185 Asylum Street, Hartford, CT 06103-3495 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))


(54) **Title:** SYSTEMS AND METHODS FOR SERVER FAILOVER AND LOAD BALANCING

FIG. 1

(57) **Abstract:** Systems and methods for server failover and/or load balancing are provided herein. Systems for server failover and load balancing may include a computer system in electronic communication over a network with one or more client applications, the computer system including a plurality of servers, and an engine stored on and executed by a client, the engine configured to allow one or more clients to select a target server among the plurality of servers using a client application identifier.

SYSTEMS AND METHODS FOR SERVER FAILOVER AND LOAD BALANCING

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims benefit under 35 U.S.C. § 119(e) to United States Provisional Patent Application No. 62/179,969 filed May 26, 2015, which is incorporated herein by reference in its entirety and made a part hereof.

FIELD OF THE DISCLOSURE

[0002] The present disclosure relates generally to systems, methods and/or media for use in server failover and/or load balancing. More specifically, the present disclosure relates to systems, methods and/or media for use in server failover and/or load balancing using clients.

BACKGROUND INFORMATION

[0003] Backend servers that provide services to clients need to be available at all times. When a particular server fails, a mechanism needs to be in place to route clients to one or more alternate servers. For example, Web services typically utilize a front-end load balancer or router, which sends requests to functioning servers. While fast, the load balancer or router is itself a critical single point of failure and introduces additional costs. Alternatively, a new server could be instantiated and the DNS (domain name system) entry for the service could be changed to point to the new server. However, there are significant delays in propagating new DNS entries.

SUMMARY OF THE INVENTION

[0004] In the mechanisms described, the routing of requests and changing of pointers is performed on the backend side (e.g., on the servers themselves) without client involvement.

[0005] It has been determined that it is possible to involve and/or otherwise utilize clients in server failover and/or load balancing.

[0006] In some embodiments, a system (and/or method and/or medium) may provide load balancing, server failover, scaling, data localization and/or redundancy. The system (and/or method and/or medium) may achieve load balancing by uniformly selecting a target server over all servers based on a random client application identifier (CAI). The system (and/or method and/or medium) may provide failover by selecting appropriate alternate servers. The system (and/or method and/or medium) may achieve scalability by extending system capacity by adding

new servers and informing directory services about the new additions. The system (and/or method and/or medium) may provide data localization and redundancy by limiting data replication to the first N servers on the CAI-specific Highest Random Weight (HRW) sorted list. The system (and/or method and/or medium) may provide the above features so when two or more clients agree on the same CAI, they will always select the same server or the same subset of servers. The system (and/or method and/or medium) may provide these benefits without any centralized services, bottlenecks, and/or single points of failure on the server side. There is no need for load balancers, third party directory involvement (e.g., DNS), or system-wide data replication.

[0007] In some embodiments, a system comprises an engine executed by one or more clients to select a target server or server subset among a plurality of servers, using a client application identifier. In some embodiments, two or more clients, using the same client application identifier, select the same target server or server subset. In some embodiments, each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services. In some embodiments, the application services on each server serve only a subset of all client applications. In some embodiments, the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports. In some embodiments, the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset. In some embodiments, the service mapping subsystem is configured to map said client application identifier using Highest Random Weight. In some embodiments, the client application identifier comprises a large number. In some embodiments, the engine is executed by a first one of the one or more clients to select a target server or server subset among the plurality of servers using the client application identifier; and the engine is executed by a second one of the one or more clients to select the same target server or server subset among the plurality of servers using the client application identifier. In some embodiments, the client application identifier is a large number. In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

[0008] In some embodiments, the system comprises a computer system that is in electronic communication over a network (or otherwise coupled (via any type(s) of communication link(s))) with one or more clients, the computer system including a plurality of servers, and an engine stored on and/or executed by the one or more clients, the engine configured to permit the one or more client applications to select a target server among the plurality of servers using a client application identifier. In some embodiments, the client application identifier is a large number. In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

[0009] In some embodiments, a method comprises electronically (or otherwise) receiving a directory list with server identities for a plurality of servers, electronically (or otherwise) receiving an application specific directory list using a client application identifier to identify servers with a desired application, applying, by an engine stored on and/or executed by a client, a Highest Random Weight to the application specific directory list to electronically (or otherwise) generate a client application identifier specific sorted list of servers that provide the desired application, and attempting by the client to electronically (or otherwise) contact a server from the client application identifier specific sorted list of servers by starting at the top and working down the list. In some embodiments, the client application identifier is a large number. In some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted. In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

[0010] In some embodiments, a method comprises: executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier. In some embodiments, two or more clients, using the same client application identifier, select the same target server or server subset. In some embodiments, each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services. In some embodiments, the application services on each server serve only a subset of all client applications. In some embodiments, the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports. In some embodiments, the engine includes a service mapping subsystem configured to map each

client application identifier to the target server or server subset. In some embodiments, the service mapping subsystem is configured to map said client application identifier using Highest Random Weight. In some embodiments, the client application identifier comprises a large number. In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: executing, by a first one of the one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier; and executing, by a second one of the one or more clients, the engine to select the target server or server subset among the plurality of servers, using the client application identifier. In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: receiving, by an engine executed by a client, information indicative of identifies of servers that provide a desired application; and generating, by the engine executed by the client, based at least in part on the information and a client application identifier, a client application identifier specific sorted list of servers that provide the desired application. In some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted. In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

[0011] In some embodiments, a non-transitory computer-readable medium has computer-readable instructions stored thereon that, if executed by a computer system, result in a method comprising: executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier. In some embodiments, two or more clients, using the same client application identifier, select the same target server or server subset. In some embodiments, each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services. In some embodiments, the application services on each server serve only a subset of all client applications. In some embodiments, the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports. In some embodiments, the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset. In some embodiments, the service mapping subsystem is configured to map said

client application identifier using Highest Random Weight. In some embodiments, the client application identifier comprises a large number. In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: executing, by a first one of the one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier; and executing, by a second one of the one or more clients, the engine to select the target server or server subset among the plurality of servers, using the client application identifier. In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: receiving, by an engine executed by a client, information indicative of identifies of servers that provide a desired application; and generating, by the engine executed by the client, based at least in part on the information and a client application identifier, a client application identifier specific sorted list of servers that provide the desired application. In some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted. In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The foregoing features of the disclosure will be apparent from the following Detailed Description, taken in connection with the accompanying drawings, in which:

[0013]

[0014] FIG. 1 is a diagram showing a server failover and load balancing system, in accordance with some embodiments;

[0015] FIG. 2 is a diagram illustrating subsystems of a failover and load balancing engine, in accordance with some embodiments;

[0016] FIG. 3 is a diagram illustrating a client communicating with a server, in accordance with some embodiments;

[0017] FIG. 4 is a flowchart of a process, in accordance with some embodiments;

[0018] FIG. 5 is a diagram showing an architecture, in accordance with some embodiments;

[0019] FIG. 6 is a graphical representation of a list (or other definition), in accordance with some embodiments;

[0020] FIG. 7 is a graphical representation of a list (or other definition), in accordance with some embodiments;

[0021] FIG. 8 is a graphical representation of a list (or other definition), in accordance with some embodiments;

[0022] FIG. 9 is a graphical representation of a list (or other definition) prior to sorting after hashing, in accordance with some embodiments;

[0023] FIG. 10 is a graphical representation of a list (or other definition), in accordance with some embodiments;

[0024] FIG. 11 is a graphical representation of a list (or other definition) prior to sorting after hashing, in accordance with some embodiments;

[0025] FIG. 12 is a graphical representation of a list (or other definition), in accordance with some embodiments; and

[0026] FIG. 13 is a flowchart of a process, in accordance with some embodiments.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0027] The present disclosure relates generally to systems and methods for server failover and/or load balancing. More specifically, the present disclosure utilizes clients, which are more numerous than servers, to address server failover and/or load balancing and/or other tasks. Accordingly, the server system is simplified as the server system is not responsible for providing uninterrupted service to clients. Instead, each client is in charge of finding a functioning server. At the same time, the server selection method guarantees that client-specific information will be available on the server chosen.

[0028] FIG. 1 is a diagram showing a server failover and load balancing system 10, in accordance with some embodiments.

[0029] Referring to FIG. 1, in accordance with some embodiments, the server failover and load balancing system 10 comprises a plurality of servers 12 (e.g., servers 12a-12n) and one or more clients 22 (e.g., clients 26a-26c). Each of the one or more clients 22 (e.g., clients 26a-26c) may include a server failover and load balancing engine (e.g., server failover and load balancing

engines 16a-16c, respectively). Each of the plurality of servers 12 may comprise any suitable type(s) of computer servers (e.g., a server with an INTEL microprocessor, multiple processors, multiple processing cores, etc.) running any suitable operating system (e.g., Windows by Microsoft, Linux, etc.). One or more of the plurality of servers 12 (e.g., servers 12a-12n) may include a database (e.g., databases 14a-14n, respectively) stored therein or operatively connected thereto. Each database may be stored on its server, or located externally therefrom (e.g., in a separate database server in communication with the server failover and load balancing system 10). As further described below, one or more of the plurality of servers 12 may also include a server failover and load balancing engine (not shown) that may be the same as and/or similar to the server failover and load balancing engines 16a-16c. Each of the plurality of servers 12 is remotely accessible such that it can communicate, through a network 20 (and/or via any other type(s) of communication link(s)), with one or more of the one or more clients 22. Each of the one or more clients 22 may include any type(s) of client device (e.g., a personal computer system 26a, a smart cellular telephone 26b, a tablet computer 26c, and/or other electronic devices). Network communication may be over the Internet using standard TCP/IP communications protocols (e.g., hypertext transfer protocol (HTTP), secure HTTP (HTTPS), file transfer protocol (FTP), electronic data interchange (EDI), etc.), through a private network connection (e.g., wide-area network (WAN) connection, emails, electronic data interchange (EDI) messages, extensible markup language (XML) messages, file transfer protocol (FTP) file transfers, etc.), or any other suitable wired or wireless electronic (or otherwise) communications format and system as should be understood by those of ordinary skill in the art.

[0030] As can be seen, the system 10 is a distributed system, which removes any single point of failure.

[0031] In accordance with some embodiments, the system 10 provides services in two basic service classes: directory services and application services. Each of the plurality of servers 12 (sometimes referred to herein as nodes) in the system 10 can provide services in one or both of such service classes (i.e., directory services and application services). In some embodiments, the application services include real time communication services (e.g., a real time communication session between two or more clients) and/or storage services (e.g., use of read/write storage by one or more clients).

[0032] In some embodiments, all directory services on all nodes contain service information about the entire system. In contrast thereto, application services (e.g., used by the client application identifier subsystem 30) on particular nodes only serve a fraction of all clients, so there is no need to distribute information about all clients to all application nodes. This reduces negative impact on system efficiency. Localization of client information on a subset of servers greatly improves network and storage performance, and enables linear capacity scaling (e.g., just by adding servers to the system).

[0033] The directory service (e.g., information directory service) provides a list of known servers, where each server is tagged or associated with all the application services it supports. In some embodiments, this list may be generated as follows as servers become “on-line” and/or otherwise available. Initially, a first server becomes “on-line” and/or otherwise available. A second server thereafter becomes “on-line” and/or otherwise available and is given the IP address of the first server. The second server subsequently sends a request to the first server for a list (or other definition) of all servers that are in the system 10 and known to the first server. After receiving a response from the first server, the second server sends a request to each additional server on the list (or other definition) for a list (or other definition) of all servers that are in the system 10 and known to such additional server. The list from the first server may initially consist of only the first server. Thus, the second server may not send any additional requests. However, the process may be repeated as each additional server becomes “on-line” and/or otherwise available and thus each additional server obtains a list (or other definition) of at least some of the servers in the system 10. Each server may also broadcast a list (or other definition) of servers known to such server, in which case, each server in the system will eventually have a list (or other definition) of each server in the system 10. In some embodiments, the list (or other definition) may also indicate the capabilities (e.g., application services) of each server on the list.

[0034] The directory services of a server may be used by one or more client computer systems to obtain a list (or other definition) of the servers that are in the system 10 and available to the one or more client computer systems. In some embodiments, a client may send a request to a server for such a list (or other definition). The client may thereafter receive a response with the list (or other definition) of servers that are in the system 10 and available to such client. In some embodiments, the client may send the request to any server in the system 10. In some

other embodiments, the client may be required to send the request to one or more designated servers (e.g., a centralized directory service). Clients may occasionally refresh this information by querying directory servers (e.g., in real time, every few hours or days, etc.).

[0035] FIG. 6 is a graphical representation of a list (or other definition) that may be received by one of the one or more clients 22, in accordance with some embodiments.

[0036] Referring to FIG. 6, the list (or other definition), which is shown in the form of a table 600 having a plurality of entries, e.g., entries 602-628, indicating all of the servers 12 that are in the system 10 and available to the client that requested the list. In the illustrated embodiment, each entry is associated with a particular one of the plurality of servers 12 available to the client, and indicates: (i) an IP address or other identifier for the server associated with the entry and (ii) one or more capabilities (e.g., application services) of such server.

[0037] For example, a first entry 602 indicates an IP address, e.g., IP address₁, associated with a first server available to the client. The first entry 602 further indicates the capabilities of such first server, e.g., real time communication. A third entry 606 indicates an IP address, e.g., IP address₃, associated with a third server available to the client and further indicates the capabilities of such third server, e.g., real time communication and storage. A ninth entry 618 indicates an IP address, e.g., IP address₉, associated with a ninth server available to the client and further indicates the capabilities of such ninth server, e.g., storage. In some embodiments, each IP address may comprise a 32 bit IP address.

[0038] In some embodiments, a client may receive and/or generate a separate list for each type of server capability (e.g., application service).

[0039] FIG. 7 is a graphical representation of a list (shown in the form of a table 700) of the servers 12 that are available to the client and have the capability of real time communication.

[0040] FIG. 8 is a graphical representation of a list (shown in the form of a table 800) of the servers 12 that are available to the client and have the capability of storage.

[0041] FIG. 2 is a diagram illustrating subsystems of a server failover and load balancing engine, e.g., one of server failover and load balancing engines 16a-16c, in accordance with some embodiments.

[0042] Referring to FIG. 2, in accordance with some embodiments, each server failover and load balancing engine, e.g. server failover and load balancing engines 16a-16c, may include a CAI subsystem 30 and a service mapping subsystem 32.

[0043] The CAI subsystem 30 facilitates failover, scaling, and/or locality of the client information. The CAI subsystem 30 associates a client application identifier (CAI) with a particular application of a particular client or group of particular clients. For example, the CAI subsystem 30 can include a real time communication session between two or more clients, use of read/write storage by one or more clients, streaming audio/video content from one client to a group of clients, application servicing a group of clients (e.g., payroll), etc. The CAI subsystem 30 defines a partition of the server system used to service a particular set of clients. This limits data and communication to the subset of the entire system, and enables scaling. The redundancy requirements are defined for each particular case within the partition (e.g., subset). If state (e.g., data) retention is involved, then the redundancy factor (e.g., two or greater) determines the replication within the subset, and subset size. If there is only a real-time component (e.g., for real time communication or live video streaming), the entire system can be involved so the redundancy factor is equal to the number of servers in the entire system. In all these cases, two or more clients that have agreed on the same CAI will always select the same server or the same subset of servers. One possible embodiment of the CAI is a large number (e.g., 128, 512, or more bits), kept secret from outsiders. It can be self-assigned (e.g., by insiders by a random choice), which can practically eliminate all collisions due to the large quantity of available possible numbers. Other embodiments are also envisioned (e.g., centralized or semi-centralized assignments). It should be understood by those of ordinary skill in the art that any suitable assignment system, method or protocol may be used.

[0044] The service mapping subsystem 32 maps each CAI to a particular server or a server subset for proper functioning of the system. Two or more clients that have agreed on the same CAI will always select the same server or the same subset of servers. This mapping works in the absence of a real-time directory service. Other than a recent directory map of the system, no other external information is required. One possible mapping method that could be used with the system is Highest Random Weight (HRW). See, e.g., https://en.wikipedia.org/wiki/Rendezvous_hashing, the entirety of which is incorporated herein by reference.

[0045] FIG. 3 is a diagram showing one of the one or more clients 22 communicating with one of the plurality of servers 12 using a list of server IP addresses 136, in accordance with some embodiments.

[0046] Referring to FIG. 3, in accordance with some embodiments, the system 10 may generate and/or otherwise provide a directory list of servers that are available to the client (with server identities represented by IP addresses). The server failover and load balancing engine in the client 22 may combine the CAI with IP addresses of servers are available to the client and provide the desired application (e.g., real time communication and/or other function) and may perform the HRW method on this list of combinations, to generate a CAI-specific sorted list of servers 136 that provide the desired application (e.g., real time communication and/or other function). In other words, the engine may sort the list by HRW (CAI, IP).

[0047] The client 22 may select a server (e.g., server 12a) from the CAI-specific sorted list of servers 136 by starting at the top of the list, and if the server at the top of the list is unavailable, the next N-1 servers 12b-12n may be tried in succession, where N is the redundancy factor. Eventually, the client 134 selects an available server (e.g., server 12d) and establishes connection with that server (e.g., server 12d).

[0048] For example, if two clients wish to establish a communication link between them, mediated by a “communication application,” the clients agree on a CAI, and run the above-described HRW algorithm. Because the directory map of both clients is identical, each of the clients will come up with the same sorted list, and will pick the same working server from that list. As there is no state in this case, all servers providing this “communication application” can be on the list, providing high redundancy.

[0049] FIG. 9 is a graphical representation of a list (or other definition) prior to sorting after hashing, in accordance with some embodiments.

[0050] Referring to FIG. 9, the list (or other definition) is shown in the form of a table 900 having a plurality of entries, e.g., entries 902-916. In the illustrated embodiment, each entry is associated with a particular one of the plurality of servers 12 and indicates: (i) the IP address or other identifier for the server associated with the entry, a combination of the CAI and the IP address of the server (e.g., concatenation of IP address and CAI) and (iii) a hash out value that results from hashing the combination of the CAI and the IP address of the server.

[0051] For example, a first entry 902 indicates an IP address, e.g., IP address₁, associated with a first server, a combination (e.g., IP address₁ & CAI) of such IP address and the CAI, and a hash out value, e.g., hash out₁, that results from hashing the combination. In some embodiments, each IP address comprises 32 bit IP address, the CAI comprises 128 bits and the combination of

the IP address and the CAI comprises 160 bits. In some embodiments, the agreed on CAI comprises a 128 bit identifier representing a telephone number.

[0052] FIG. 10 is a graphical representation of the list (or other definition) of IP addresses shown in table 900 after sorting based on the hash out values in the table 900, in accordance with some embodiments.

[0053] Referring to FIG. 10, the sorted list (or other definition) is shown in the form of a table 1000 having a plurality of entries, e.g., 1002-1016. For purposes of illustration, it has been assumed that the hash out values in the table 900 define the following descending sequence of hash out values (greatest to smallest): hash out₇, hash out₃, hash out₁, hash out₂, hash out₆, hash out₅, hash out₄, hash out₈.

[0054] As stated above, the sorted list is not limited to the forms described above. In accordance with some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted.

[0055] The above communication example is just one non-limiting example, and many other applications are possible as should be understood by those of ordinary skill in the art.

[0056] In another example, a group of clients belonging to the same business wants to access payroll. They all agree on a CAI, and select a functioning server, which has access to payroll data related to this group (e.g., defined by CAI). The application services in this case have state (e.g., payroll information), and a redundancy requirement with factor N, usually greater than 1. To achieve this, in some embodiments, each payroll application service (and/or other portion of the functioning server), when it receives a request/session from a particular client, also receives CAI, and calculates (or otherwise generates) the HRW sorted list itself. Then it finds the first N servers on the list (e.g., one of them is itself), and replicates all payroll information for this CAI to the other N-1 servers on the list. In some embodiments, the functioning server receives the HRW sorted list (or a portion thereof) from the requesting client and does not generate the HRW sorted list itself. In some embodiments, the functioning server does not replicate the information to the other N-1 servers on the list. In the latter embodiments, the requesting client may replicate the information to the other N-1 servers.

[0057] FIG. 11 is a graphical representation of a list (or other definition) prior to sorting after hashing, in accordance with some embodiments.

[0058] Referring to FIG. 11, the list (or other definition) is shown in the form of a table 1100 having a plurality of entries, e.g., entries 1102-1116. In the illustrated embodiment, each entry is associated with a particular one of the plurality of servers 12 and indicates: (i) the IP address or other identifier for the server associated with the entry, a combination of the CAI and the IP address of the server (e.g., concatenation of IP address and CAI) and (iii) a hash out value that results from hashing the combination of the CAI and the IP address of the server.

[0059] For example, a first entry 1102 indicates an IP address, e.g., IP address₃, associated with a third server, a combination (e.g., IP address₃ & CAI) of such IP address and the CAI, and a hash out value, e.g., hash out₃, that results from hashing the combination. In some embodiments, each IP address comprises 32 bit IP address, the CAI comprises 104 bits and the combination of the IP address and the CAI comprises 136 bits.

[0060] FIG. 12 is a graphical representation of the list (or other definition) of IP addresses shown in table 1100 after sorting based on the hash out values in the table 1100, in accordance with some embodiments.

[0061] Referring to FIG. 12, the sorted list (or other definition) is shown in the form of a table 1200 having a plurality of entries, e.g., 1202-1216. For purposes of illustration, it has been assumed that the hash out values in the table 1100 define the following descending sequence of hash out values (greatest to smallest): hash out₇, hash out₃, hash out₉, hash out₁₀, hash out₆, hash out₅, hash out₄, hash out₈.

[0062] In some embodiments, the client sends the selected server information to identify a record (of any size) within the selected server at which to store the data (if the data is being stored) and/or from which to retrieve the data (if the data is being retrieved). In some embodiments, the information comprises 128 bits. In some embodiments, the information may comprise the CAI, which the selected server may use as a key to identify a partition within the selected server, and a record number that identifies a record within the identified partition. In some embodiments, the CAI used to identify the partition may comprise 104 bits and the record number used to identify the record within the storage partition may comprise 24 bits.

[0063] As stated above, the sorted list is not limited to the forms described above. In accordance with some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted.

[0064] The above storage example is just one more non-limiting example, and many other applications are possible as should be understood by those of ordinary skill in the art.

[0065] FIG. 4 is a flowchart of the process 150, in accordance with some embodiments. In step 152, the system 10 electronically (or otherwise) generates a directory list with server identities represented by IP addresses. In step 154, the system electronically (or otherwise) generates an application specific directory list using a client application identifier to identify servers with a desired application. In step 156, the client electronically (or otherwise) applies the HRW method on the application specific directory list to electronically (or otherwise) generate a CAI-specific sorted list of servers that provide the desired application. In step 158, one or more clients attempt to contact the first server listed in the CAI-specific sorted list of servers. In step 160, a determination is made as to whether an answer was received from the server selected. If not, then in step 162, the client attempts to contact the next server listed in the CAI-specific sorted list of servers. If so, then in step 164, an electrical or data communication is established between the one or more clients and the server selected.

[0066] FIG. 13 is a flow chart of a process 1300, in accordance with some embodiments.

[0067] In accordance with some embodiments, the method may be used in server failover and/or load balancing.

[0068] The method 1300 is not limited to the order shown in the flow chart. Rather, embodiments of the method 1300 may be performed in any order that is practicable. For that matter, unless stated otherwise, any method disclosed herein may be performed in any order that is practicable.

[0069] Unless stated otherwise, the method 1300 (and/or any other method disclosed herein) may be performed by in any manner. In some embodiments, the method, (and/or any other method disclosed herein) or one or more portions thereof, may be performed by one or more portions of the system 10 and/or any other processing system.

[0070] In some embodiments, a non-transitory computer readable medium may have instructions stored thereon, which if executed by a machine result in performance of the method 1300 (and/or any other method disclosed herein) or one or more portions thereof.

[0071] Referring to FIG. 13, in accordance with some embodiments, at 1302, the method may include executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier.

[0072] In some embodiments, two or more clients, using the same client application identifier, select the same target server or server subset.

[0073] In some embodiments, the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset.

[0074] In some embodiments, the service mapping subsystem is configured to map said client application identifier using Highest Random Weight.

[0075] In some embodiments, the client application identifier comprises a large number.

[0076] In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: executing, by a first one of the one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier; and executing, by a second one of the one or more clients, the engine to select the target server or server subset among the plurality of servers, using the client application identifier.

[0077] In some embodiments, executing an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises: receiving, by an engine executed by a client, information indicative of identifies of servers that provide a desired application; and generating, by the engine executed by the client, based at least in part on the information and a client application identifier, a client application identifier specific sorted list of servers that provide the desired application.

[0078] In some embodiments, the sorted list may comprise any type(s) of information, in any form(s), that defines an order in which to attempt to contact servers until at least one functioning one of the servers is contacted.

[0079] In some embodiments, the large number comprises at least 128 bits. In some embodiments, the large number comprises at least 512 bits.

[0080] FIG. 5 is a diagram showing an architecture 200, in accordance with some embodiments. In some embodiments, one or more of the systems and/or devices (and/or portion(s) thereof) disclosed herein may have an architecture that is the same as and/or similar to one or more portions of the architecture 200. In some embodiments, one or more of the plurality of servers 12 in the server failover and load balancing system 10 and/or one or more of the one or more clients 22 in the server failover and load balancing system 10 have an architecture that is the same as and/or similar to one or more portions of the architecture 200.

[0081] In some embodiments, one or more of the methods (or portion(s) thereof) disclosed herein may be performed by a system, apparatus and/or device having an architecture that is the same as or similar to the architecture 200 (or portion(s) thereof).

[0082] The architecture may be implemented as a distributed architecture or a non-distributed architecture. A distributed architecture may be a completely distributed architecture or a partly distributed-partly non distributed architecture.

[0083] Referring to FIG. 5, in accordance with some embodiments, the architecture 200 includes a processing system 202 that may include one or more of a storage device 204, a network interface 208, a communications bus 210, a central processing unit (CPU) (microprocessor) 212, a random access memory (RAM) 214, and one or more input devices 216, such as a keyboard, mouse, etc. The processing system 202 may also include a display (e.g., liquid crystal display (LCD), cathode ray tube (CRT), etc.). In some embodiments, the storage device 204 may comprise any suitable, computer-readable storage medium such as disk, non-volatile memory (e.g., read-only memory (ROM), erasable programmable ROM (EPROM), electrically-erasable programmable ROM (EEPROM), flash memory, field-programmable gate array (FPGA), etc.). In some embodiments, the processing system 202 may comprise a networked computer system, a personal computer, a smart phone, a tablet computer etc. In some embodiments, a server failover and load balancing engine 16 may be embodied as computer-readable program code stored on the storage device 204 and executed by the CPU 212 using any suitable, high or low level computing language, such as, for example, Python, Java, C, C++, C#, .NET, MATLAB, etc. In some embodiments, the network interface 208 may include an Ethernet network interface device, a wireless network interface device operable with one or more suitable wireless protocol, or any other suitable device that permits the server 202 to communicate via the network. In some embodiments, the CPU 212 may include any suitable single or multiple-core microprocessor of any suitable architecture that is capable of implementing and running the server failover and load balancing engine 16. In some embodiments, the random access memory 214 may include any suitable, high-speed, random access memory typical of most modern computers, such as dynamic RAM (DRAM), etc.

[0084] Unless stated otherwise, any process (sometimes referred to herein as a method) disclosed herein is not limited to an order shown in a flow chart. Rather, embodiments may be performed in any order that is practicable.

[0085] Unless stated otherwise, any method disclosed herein may be performed in any manner. In some embodiments, a method (or one or more portions thereof) may be performed by one or more portions of the system 10.

[0086] In some embodiments, a non-transitory computer readable medium may have instructions stored thereon, which if executed by a machine result in performance of any one or more methods (or one or more portions thereof) disclosed herein.

[0087] Unless stated otherwise, a processor may comprise any type of processor. For example, a processor may be programmable or non-programmable, general purpose or special purpose, dedicated or non-dedicated, distributed or non-distributed, shared or not shared, and/or any combination thereof. A processor may include, but is not limited to, hardware, software, firmware, and/or any combination thereof. Hardware may include, but is not limited to off the shelf integrated circuits, custom integrated circuits and/or any combination thereof. In some embodiments, a processor comprises a microprocessor. Software may include, but is not limited to, instructions that are storable and/or stored on a computer readable medium, such as, for example, magnetic or optical disk, magnetic or optical tape, CD-ROM, DVD, RAM, EPROM, ROM or other semiconductor memory. A processor may employ continuous signals, periodically sampled signals, and/or any combination thereof. If a processor is distributed, two or more portions of the control/storage circuitry may communicate with one another through a communication link.

[0088] Unless stated otherwise, a processing system is any type of system that includes at least one processor.

[0089] Unless stated otherwise, a communication link may be any type of communication link, for example, but not limited to, wired (e.g., conductors, fiber optic cables) or wireless (e.g., acoustic links, electromagnetic links or any combination thereof including, for example, but not limited to microwave links, satellite links, infrared links), and/or combinations thereof, each of which may be public or private, dedicated and/or shared (e.g., a network). A communication link may or may not be a permanent communication link. A communication link may support any type of information in any form, for example, but not limited to, analog and/or digital (e.g., a sequence of binary values, i.e. a bit string) signal(s) in serial and/or in parallel form. The information may or may not be divided into blocks. If divided into blocks, the amount of information in a block may be predetermined or determined dynamically, and/or may be fixed

(e.g., uniform) or variable. A communication link may employ a protocol or combination of protocols including, for example, but not limited to the Internet Protocol.

[0090] Unless stated otherwise, the term “receiving” means receiving in any manner(s) from any source(s) internal and/or external.

[0091] Unless stated otherwise the term "define" means "define directly and/or indirectly".

[0092] Unless stated otherwise, terms such as, for example, "in response to" and "based on" mean "in response at least to" and "based at least on", respectively, so as not to preclude being responsive to and/or based on, more than one thing.

[0093] Unless stated otherwise, terms such as, for example, “coupled to”, “send to” and “receive from” mean “coupled directly and/or indirectly to”, “send directly and/or indirectly to” and “receive directly and/or indirectly from”, respectively.

[0094] Unless stated otherwise, terms such as, for example, "comprises", "has", "includes", and all forms thereof, are considered open-ended, so as not to preclude additional elements and/or features. In addition, unless stated otherwise, terms such as, for example, "a", "one", "first", are considered open-ended, and do not mean “only a”, “only one” and “only a first”, respectively. Moreover, unless stated otherwise, the term “first” does not, by itself, require that there also be a “second”.

[0095]

Having thus described the system and method in detail, it is to be understood that the foregoing description is not intended to limit the spirit or scope thereof. It will be understood by those of ordinary skill in the pertinent art that the embodiments of the present disclosure described herein are merely exemplary and that a person skilled in the art, based on the teachings herein, may make any numerous variations and modifications without departing from the spirit and/or scope of the disclosure. All such variations and modifications, including those discussed above, are intended to be included within the scope of the disclosure. By way of example only, the disclosure contemplates, but is not limited to, embodiments having any one or more of the features set forth in the above description in any combination with each other. By way of another example only, the disclosure contemplates, but is not limited to, embodiments that do not include any one or more of the features set forth in the above description. Accordingly, the description of embodiments herein is to be taken in an illustrative as opposed to a limiting sense,

and neither the invention nor the disclosure is limited to the particular embodiments or combinations of features described herein.

What is claimed is:

1. A system comprising:
an engine executed by one or more clients to select a target server or server subset among a plurality of servers, using a client application identifier.
2. The system of claim 1, wherein two or more clients, using the same client application identifier, select the same target server or server subset.
3. The system of claim 1, wherein each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services.
4. The system of claim 3, wherein the application services on each server serve only a subset of all client applications.
5. The system of claim 3, wherein the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports.
6. The system of claim 1, wherein the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset.
7. The system of claim 6, wherein the service mapping subsystem is configured to map said client application identifier using Highest Random Weight.
8. The system of any of claims 1 to 7, wherein the client application identifier comprises a large number.
9. The system of any of claim 1 to 8, wherein the engine is executed by a first one of the one or more clients to select a target server or server subset among the plurality of servers using the client application identifier; and,
wherein the engine is executed by a second one of the one or more clients to select the same target server or server subset among the plurality of servers using the client application identifier.

10. A method comprising:
executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier.
11. The method of claim 10, wherein two or more clients, using the same client application identifier, select the same target server or server subset.
12. The method of claim 10, wherein each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services.
13. The method of claim 12, wherein the application services on each server serve only a subset of all client applications.
14. The method of claim 12, wherein the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports.
15. The method of claim 10, wherein the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset.
16. The method of claim 15, wherein the service mapping subsystem is configured to map said client application identifier using Highest Random Weight.
17. The method of any of claims 10 to 16, wherein the client application identifier comprises a large number.
18. The method of any of claims 10 to 17, wherein the executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises:
executing, by a first one of the one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier;
and

executing, by a second one of the one or more clients, the engine to select the target server or server subset among the plurality of servers, using the client application identifier.

19. A non-transitory computer-readable medium having computer-readable instructions stored thereon that, if executed by a computer system, result in a method comprising:
executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier.

20. The non-transitory computer-readable medium of claim 19, wherein two or more clients, using the same client application identifier, select the same target server or server subset.

21. The non-transitory computer-readable medium of claim 19, wherein each of the plurality of servers is coupled to the one or more clients and configured to provide one or more service classes, the services classes including application services and directory services.

22. The non-transitory computer-readable medium of claim 21, wherein the application services on each server serve only a subset of all client applications.

23. The non-transitory computer-readable medium of claim 21, wherein the directory services of each server are configured to provide a list of servers, each server tagged with the application services that server supports.

24. The non-transitory computer-readable medium of claim 19, wherein the engine includes a service mapping subsystem configured to map each client application identifier to the target server or server subset.

25. The non-transitory computer-readable medium of claim 24, wherein the service mapping subsystem is configured to map said client application identifier using Highest Random Weight.

26. The non-transitory computer-readable medium of any of claims 19 to 25, wherein the client application identifier comprises a large number.

27. The non-transitory computer-readable medium of any of claims 19 to 26, wherein the executing, by one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier comprises:

executing, by a first one of the one or more clients, an engine to select a target server or server subset among a plurality of servers, using a client application identifier;
and

executing, by a second one of the one or more clients, the engine to select the target server or server subset among the plurality of servers, using the client application identifier.

1/13




FIG. 1

2/13




FIG. 2




FIG. 3

4/13




FIG. 4

5/13





FIG. 5

6/13



The diagram shows a table with three columns: ENTRY, SERVER IP ADDRESS, and SERVER CAPABILITIES. The table is labeled 600. The first row is labeled 602. The second row is labeled 604. The third row is labeled 606. The fourth row is labeled 608. The fifth row is labeled 610. The sixth row is labeled 612. The seventh row is labeled 614. The eighth row is labeled 616. The ninth row is labeled 618. The tenth row is labeled 620. The eleventh row is labeled 622. The twelfth row is labeled 624. The thirteenth row is labeled 626. The fourteenth row is labeled 628.

ENTRY	SERVER IP ADDRESS	SERVER CAPABILITIES
1	IP ADDRESS ₁	COMMUNICATION
2	IP ADDRESS ₂	COMMUNICATION
3	IP ADDRESS ₃	COMMUNICATION STORAGE
4	IP ADDRESS ₄	COMMUNICATION STORAGE
5	IP ADDRESS ₅	COMMUNICATION STORAGE
6	IP ADDRESS ₆	COMMUNICATION STORAGE
7	IP ADDRESS ₇	COMMUNICATION STORAGE
8	IP ADDRESS ₈	COMMUNICATION STORAGE
9	IP ADDRESS ₉	STORAGE
10	IP ADDRESS ₁₀	STORAGE
⋮	⋮	⋮
n-1	IP ADDRESS _{n-1}	⋮
n	IP ADDRESS _n	⋮

FIG. 6

700

ENTRY	SERVER IP ADDRESS	SERVER CAPABILITIES
1	IP ADDRESS ₁	COMMUNICATION
2	IP ADDRESS ₂	COMMUNICATION
3	IP ADDRESS ₃	COMMUNICATION STORAGE
4	IP ADDRESS ₄	COMMUNICATION STORAGE
5	IP ADDRESS ₅	COMMUNICATION STORAGE
6	IP ADDRESS ₆	COMMUNICATION STORAGE
7	IP ADDRESS ₇	COMMUNICATION STORAGE
8	IP ADDRESS ₈	COMMUNICATION STORAGE

FIG. 7

800

ENTRY	SERVER IP ADDRESS	SERVER CAPABILITIES
1	IP ADDRESS ₃	COMMUNICATION STORAGE
2	IP ADDRESS ₄	COMMUNICATION STORAGE
3	IP ADDRESS ₅	COMMUNICATION STORAGE
4	IP ADDRESS ₆	COMMUNICATION STORAGE
5	IP ADDRESS ₇	COMMUNICATION STORAGE
6	IP ADDRESS ₈	COMMUNICATION STORAGE
7	IP ADDRESS ₉	STORAGE
8	IP ADDRESS ₁₀	STORAGE

FIG. 8

900

ENTRY	SERVER IP ADDRESS	IP ADDRESS COMBINED WITH CLIENT APPLICATION IDENTIFIER (CAI)	HASH OUT
1	IP ADDRESS ₁	IP ADDR ₁ & CAI	HASH OUT ₁
2	IP ADDRESS ₂	IP ADDR ₂ & CAI	HASH OUT ₂
3	IP ADDRESS ₃	IP ADDR ₃ & CAI	HASH OUT ₃
4	IP ADDRESS ₄	IP ADDR ₄ & CAI	HASH OUT ₄
5	IP ADDRESS ₅	IP ADDR ₅ & CAI	HASH OUT ₅
6	IP ADDRESS ₆	IP ADDR ₆ & CAI	HASH OUT ₆
7	IP ADDRESS ₇	IP ADDR ₇ & CAI	HASH OUT ₇
8	IP ADDRESS ₈	IP ADDR ₈ & CAI	HASH OUT ₈

902

904

906

908

910


912

914

916

FIG. 9

10/13



The diagram shows a table with 8 rows and 2 columns. The first column is labeled 'ENTRY' and the second column is labeled 'SERVER IP ADDRESS'. The rows are numbered 1 through 8. The server IP addresses are listed as IP ADDRESS₇, IP ADDRESS₃, IP ADDRESS₁, IP ADDRESS₂, IP ADDRESS₆, IP ADDRESS₅, IP ADDRESS₄, and IP ADDRESS₈. Arrows point from the labels 1000, 1002, 1004, 1006, 1008, 1010, 1012, 1014, and 1016 to the corresponding rows of the table.

ENTRY	SERVER IP ADDRESS
1	IP ADDRESS ₇
2	IP ADDRESS ₃
3	IP ADDRESS ₁
4	IP ADDRESS ₂
5	IP ADDRESS ₆
6	IP ADDRESS ₅
7	IP ADDRESS ₄
8	IP ADDRESS ₈

HASH OUT₇ > HASH OUT₃ > HASH OUT₁ > HASH OUT₂
HASH OUT₆ > HASH OUT₅ > HASH OUT₄ > HASH OUT₈

FIG. 10

11/13

1100

ENTRY	SERVER IP ADDRESS	IP ADDRESS COMBINED WITH CLIENT APPLICATION IDENTIFIER (CAI)	HASH OUT
1	IP ADDRESS ₃	IP ADDR ₃ & CAI	HASH OUT ₃
2	IP ADDRESS ₄	IP ADDR ₄ & CAI	HASH OUT ₄
3	IP ADDRESS ₅	IP ADDR ₅ & CAI	HASH OUT ₅
4	IP ADDRESS ₆	IP ADDR ₆ & CAI	HASH OUT ₆
5	IP ADDRESS ₇	IP ADDR ₇ & CAI	HASH OUT ₇
6	IP ADDRESS ₈	IP ADDR ₈ & CAI	HASH OUT ₈
7	IP ADDRESS ₉	IP ADDR ₉ & CAI	HASH OUT ₉
8	IP ADDRESS ₁₀	IP ADDR ₁₀ & CAI	HASH OUT ₁₀

1102
1104
1106
1108
1110
1112
1114
1116


FIG. 11

12/13

ENTRY	SERVER IP ADDRESS
1	IP ADDRESS ₇
2	IP ADDRESS ₃
3	IP ADDRESS ₉
4	IP ADDRESS ₁₀
5	IP ADDRESS ₆
6	IP ADDRESS ₅
7	IP ADDRESS ₄
8	IP ADDRESS ₈

HASH OUT₇ > HASH OUT₃ > HASH OUT₉ > HASH OUT₁₀ >
HASH OUT₆ > HASH OUT₅ > HASH OUT₄ > HASH OUT₈

FIG. 12

13/13**FIG. 13**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US16/34493

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - H04L 29/08; G06F 9/50, 15/173 (2016.01)

CPC - G06F 9/5055; H04L 67/1002, 67/1008, 67/42

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8): G06F 9/46, 9/50, 15/173; H04L 29/08 (2016.01)

CPC: G06F 9/46, 9/505, 9/5055, 15/173; H04L 67/1002, 67/1008, 67/42

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, IN, CA, Other Countries (INPADOC), RU, AT, CH, TH, BR, PH); EBSCO; IEEE/IEEEExplore; Google/Google Scholar; Keywords: failover, load balance, server, host, application, services, server, client, device, identifier, engine, proxy, select, direct, subset

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ----- Y	US 2010/0333104 A1 (UNNIKRISHNAN, U et al.) 30 December 2010; Figure 1; paragraphs [0028], [0029], [0066]; claim 5.	1-3, 8/1-8/3, 10-12, 17/10-17/12, 19-21, 26/19-26/21 ----- 4-7, 8/4-8/7, 13-16, 17/13-17/16, 22-25, 26/22-26/25
Y	US 8,902,787 B2 (QIU, C. et al.) 02 December 2014; column 5, lines 3-10; column 7, lines 25-28; column 8, lines 56-67; claim 1.	4-5, 8/4-8/5, 13-14, 17/13-17/14, 22-23, 26/22-26/23
Y	US 8,150,957 B1 (MASTERS, R. et al.) 03 April 2012; column 14, line 50 – column 15, line 5.	6-7, 8/6-8/7, 15-16, 17/15-17/16, 24-25, 26/24-26/25
Y	US 2015/0063364 A1 (NICIRA, INC.) 05 March 2015; paragraphs [0183], [0188].	7, 8/7, 16, 17/16, 25, 26/25

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

“A” document defining the general state of the art which is not considered to be of particular relevance

“E” earlier application or patent but published on or after the international filing date

“L” document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

“O” document referring to an oral disclosure, use, exhibition or other means

“P” document published prior to the international filing date but later than the priority date claimed

“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

“&” document member of the same patent family

Date of the actual completion of the international search

27 July 2016 (27.07.2016)

Date of mailing of the international search report

19 AUG 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US16/34493

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☒ Claims Nos.: 9, 18, 27
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.