

(19) DANMARK

(10) DK/EP 3143753 T3



(12)

Oversættelse af
europæisk patentskrift

Patent- og
Varemærkestyrelsen

(51) Int.Cl.: **H 04 L 29/08 (2006.01)**

(45) Oversættelsen bekendtgjort den: **2018-12-03**

(80) Dato for Den Europæiske Patentmyndigheds
bekendtgørelse om meddelelse af patentet: **2018-08-29**

(86) Europæisk ansøgning nr.: **15726434.2**

(86) Europæisk indleveringsdag: **2015-05-11**

(87) Den europæiske ansøgnings publiceringsdag: **2017-03-22**

(86) International ansøgning nr.: **US2015030235**

(87) Internationalt publikationsnr.: **WO2015175442**

(30) Prioritet: **2014-05-13 US 201461992623 P** **2014-09-24 US 201414495683**

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

(73) Patenthaver: **Google LLC, 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA**

(72) Opfinder: **EISENBUD, Daniel Eugene, 1600 Amphitheatre Parkway, Mountain View, California 94043, USA**
NEWTON, Simon Jeffrey, 1600 Amphitheatre Parkway, Mountain View, California 94043, USA

(74) Fuldmægtig i Danmark: **AWA Denmark A/S, Strandgade 56, 1401 København K, Danmark**

(54) Benævnelse: **FREMGANGSMÅDE OG SYSTEM TIL BELASTNINGSBALANCERING AF ANYCAST-DATATRAFIK**

(56) Fremdragne publikationer:
US-A1- 2006 212 597
US-A1- 2010 302 940
US-A1- 2011 145 390
US-B1- 7 355 977

DK/EP 3143753 T3

DESCRIPTION

FIELD OF THE DISCLOSURE

[0001] The present disclosure relates generally to the field of data traffic load balancing.

BACKGROUND

[0002] Much of the world's Internet-accessible content is provided by servers hosted in large data centers. Such data centers are typically distributed across multiple geographical locations, and serve end-users globally. A typical data center houses thousands of servers hosting multiple instances of software applications associated with different services provided to end-users. When an end-user makes a request for Internet content associated with, for example, a web site, social media service, streaming media content item, gaming service, or any other online service, the request is sent to an application instance associated with the service hosted by a data center to serve request.

[0003] Patent Application US 2010/302940 discloses a load balancing system which includes two layers (two levels) of load balancers. The first level of dynamic load balancers DLB distribute / forward the packets to local load balancers LLB. Packets of a particular service are directed (destination address) to a single Virtual IP and are then forwarded by routers to the DLB (first level). The DLB use a hashing function on the packet header fields to select the target device (second level LLB or application servers).

[0004] Patent Application US 7355977 discloses an algorithm to distribute received packets to targets / servers (load balancing). A hash on header fields is an index to a mapping table where each slot contains the identifier of a target (e.g. IP). The number of occurrences of a target identifier in the table is proportional to the weight (load ration) of this target, i.e. number of slots for each server is calculated as the proportion of the load share.

[0005] Patent Application US 2011/145390 discloses a load balancing method wherein an virtual Internet protocol address for forwarding client packets by routers to a set of load balancing devices is envisaged. The load balancers use hashing of packet header fields mapped to target servers for session persistence and an uneven (weighted) load distribution to available servers. Routers identify a next hop and forward packets based on destination information (e.g., a destination IP address and/ or port) included in a header of a received packet.

SUMMARY

[0006] According to one aspect of the disclosure, a method for load balancing data traffic addressed to an anycast address includes maintaining, by each load balancing (LB) device of a first set of LB devices, a first data structure including entries associated with application instances in a group of application instances served by the LB device of the first set of LB devices. The frequency with which each served application instance is included in the first data structure is indicative of a weight value associated with a capacity of the corresponding served application instance. Upon an LB device in the first set of LB devices receiving a data packet originally addressed to an anycast address, the LB device generates a first hash value based on a first set of header fields of the received data packet. The LB of the first set then uses the first data structure to identify a virtual Internet protocol (IP) address of one of the served application instances based on the generated first hash value, and forwards the data packet to the identified application instance. The method also includes maintaining, by a LB device of a second set of LB devices in the first set, the frequency with which each LB device in the first set of LB devices is included in the second data structure is indicative of a weight value associated with the corresponding LB device of the first set. Upon an LB device of the second set of LB devices receiving a data packet originally addressed to the anycast address, the LB generates a second hash value based on a second set of header fields of the received data packet, and identifies a LB device of the first set using the second data structure, based on the generated second hash value. The LB of the second set then forwards the data packet to the identified LB device of the first set of LB devices.

[0007] According to another aspect of the disclosure, a system for load balancing anycast traffic in a communications network includes a first set of load balancing (LB) devices. Each LB device of the first set of LB devices is configured to maintain a first data structure including entries associated with application instances in a group of application instances served by the LB device of the first set of LB devices. The frequency with which each served application instance is included in the first data structure is indicative of a weight value associated with a corresponding served application instance. Upon an LB device of the first set of LB devices receiving a data packet, received at the system addressed to an anycast address, the LB device generates a first hash value based on one or more first header fields of the received data packet, and uses the first data structure to identify a virtual Internet protocol (IP) address of one of the served application instances based on the generated first hash value. The LB device of the first set then forwards the data packet to the identified application instance. The system also includes a second set of load balancing LB devices. Each LB device of the second set of LB devices is configured to maintain a second data structure including entries associated with respective LB devices in the first set. The frequency with which each LB device in the first set of LB devices is included in the second data structure is indicative of a weight value associated with the corresponding LB device of the first set. Upon an LB device in the second set of LB devices receiving a data packet, received at the system addressed to the anycast address, the LB device generates a second hash value based on one or more second header fields of the received data packet, and identifies a LB device of first set of LB devices using the second data structure, based on the generated second hash value. The LB device of the second set then and forwards the data packet to the identified LB device of the first set. The

system further includes an anycast node associated with the anycast address configured to, upon receiving a data packet addressed to the anycast address, forward the received data packet to a LB device in the second set of LB devices.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The above and related objects, features, and advantages of the present disclosure will be more fully understood by reference to the following detailed description, when taken in conjunction with the following figures, wherein:

Figure 1 shows a block diagram of an implementation of a single-layer load balancing system for load balancing data traffic addressed to an anycast address;

Figure 2 shows a flowchart describing an implementation of a process of processing anycast data packets performed by the single-layer load balancing system ;

Figure 3 shows a block diagram of an implementation of a two-layer load balancing system for load balancing data traffic addressed to an anycast address;

Figure 4 shows a block diagram representing another implementation of the two-layer load balancing system;

Figure 5 shows a flowchart describing an implementation of a process of handling anycast data packets performed by the two-layer load balancing system;

Figure 6 shows a flowchart describing an implementation of a process for generating a data structure employed by load balancers; and

Figure 7 shows illustrations of data structures employed by systems in Figures 1 and 3.

[0009] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0010] Online services or applications usually include multiple application instances for each application or service. The multiple application instances for each application or service may reside on multiple computer servers. As such, the load associated with accessing an application or service by multiple end-users may be distributed across at least a subset of the corresponding multiple application instances. Also, distributing the multiple application instances for a given application or service across different geographical locations helps reduce latency experienced by end-users. In particular, an end-user may be served by the

closest geographical location having one or more of the corresponding multiple application instances. While serving each end-user or client by the corresponding closest geographical location reduces latency, the computer servers in each geographical location, and the application instances therein, have a finite capacity. When demand for a given application or service at a particular location exceeds the capacity of the applications instances in the same location, excess demand for the application may overflow to the next closest location.

[0011] In order to address the finite capacities of the computer servers and the application instances executed thereon, domain name system (DNS) servers have been traditionally employed in existing data centers to perform load balancing functionalities. In the following, implementations of processes, apparatuses, and systems for load balancing anycast data traffic are presented.

[0012] In anycast-based services or applications, one or more Internet protocol (IP) addresses are advertised, from multiple servers, globally using anycast. The anycast IP address for a given application or service is then used by end-users when accessing the same application or service. Data traffic associated with the anycast address, e.g., requests from end-users to access the application or service, is then load balanced across different corresponding application instances as depicted in the implementations described below. Data traffic associated with an anycast address is also referred to herein as anycast traffic or anycast data traffic. In some implementation, anycast traffic associated with stateful protocols, e.g., transport control protocol (TCP), is load balanced and served while corresponding connections are maintained even as Internet routing tables change. Furthermore, implementations described below can allow for specifying capacities or other load balancing metrics for different locations, and allow for rapid load balancing responses to changes in load, capacity, or any other load balancing metrics.

[0013] Figure 1 shows a block diagram of an implementation of a single layer load balancing system 100 for load balancing data traffic addressed to an anycast address. The system 100 includes multiple anycast redirector nodes 110a - 110c (also referred to hereinafter either individually or collectively as anycast node(s) 110) load balancer devices 120a - 120b (also referred to hereinafter either individually or collectively as load balancer(s) 120) and multiple server clusters 130a - 130c (also referred to hereinafter either individually or collectively as cluster(s) 130). Each of the clusters 130 includes a number of application instances 131.

[0014] A data packet 10 addressed to an anycast address is received at the system 100 by an anycast node 110. For example, the data packet 10 is received by the anycast node 110. The anycast nodes 110 are devices configured to receive data packets 10 addressed to an anycast address, and redirect the received data packets to corresponding load balancers 120. In some instances, an anycast node 110 redirects a received data packet 10 to a corresponding load balancer 120. In other instances, the receiving anycast node 110 redirects the received data packet 10 to a pool of load balancers 120 serving one or more respective clusters 130 of the system 100.

[0015] In some implementations, each anycast node 110 maintains, or has access to, a respective map mapping source Internet Protocol (IP) addresses to corresponding clusters 130 or load balancers 120 associated with such clusters. The map is also referred to as a source IP map. In some instances, each source IP address is mapped to the cluster 130, or corresponding load balancer(s) 120, that is closest to the data packet source associated with the same source IP address. A person skilled in the art should appreciate that the mappings between source IP addresses and clusters 130, or load balancers 120, may be defined differently, for example, based on assignments made by administrators of the system 100. Upon receiving the data packet 10 addressed to the anycast address, the anycast node 110 looks up the packet's source Internet Protocol (IP) address in the source IP map, and redirects the data packet 10 to a mapped location, *i.e.*, a location associated with the source IP address in the source IP map. The mapped location is indicative of a load balancer 120 or a pool of load balancers 120 of the system 100.

[0016] In some implementations, the anycast node 110 uses several header fields of the received data packet 10 to determine a load balancer 120, or a pool of load balancers 120, to which the received data packet 10 is to be redirected. For instance, the receiving anycast node 110 uses a destination IP address (such as a virtual IP (VIP) address) and a source port associated with the received data packet 10 to determine an anycast service group among multiple anycast service groups. In some implementations, the receiving anycast node 110 can use the destination address, the destination port, the source IP address, the source port, the protocol, or any combination thereof to determine the anycast service group. The receiving anycast node 110 may also use other information included in the data packet payload such as a connection identifier (ID) or the like. For example, two anycast service groups may be defined, one using end-point servers for backend load balancing and another employing load balancing servers for backend load balancing. End-point servers are configured to terminate a transport control protocol (TCP) connection. In other examples, different anycast service groups may be defined. In mapping the packet's destination IP address and source port to the corresponding anycast service group, the receiving anycast node 110 makes use, for example, of a second map or configuration information indicative of such mappings. The receiving anycast node 110 further employs the source IP address of the received data packet 10 and the source IP map associated with receiving anycast node 110 to determine a zone, *e.g.*, one or more clusters 130, of the system 100. In some implementations, a single source IP map is shared by all anycast nodes 110. In other implementations, different anycast nodes 110 may be associated with distinct source IP maps. In some implementations, each source IP map is associated with a respective service. The receiving anycast node 110 then maps the determined zone and the determined anycast service group to a load balancer 120 or a pool of load balancers 120 based on, for example, a third map. The receiving anycast node then redirects the received data packet 10 to the determined load balancer 120 or the determined pool of load balancers 120.

[0017] A person of ordinary skill in the art should appreciate that using a global mapping, *e.g.*, a source IP map, makes redirecting of data packets 10 by the anycast nodes 110 consistent. In other words, if the data packet is received at the anycast node 110c instead of the anycast

110b, the anycast node 110c will, using the source IP map, redirect the data packet 10 to the same load balancer that would have been selected by the anycast node 110b.

[0018] The receiving anycast node 110 can be further configured to check a connection table to determine whether the received data packet 10 corresponds to a data flow, or session, recently served by the receiving anycast node 110. If the received data packet corresponds to a data flow or session indicated in the connection table, the receiving anycast node 110 forwards the received data packet 10 to a destination associated with the corresponding data flow or session. Otherwise, the receiving anycast node 110 determines a load balancer 120, or a pool of load balancers 120, as described above, and redirects the received data packet 10 to the determined load balancer 120 or the determined pool of load balancers 120.

[0019] In some implementations, each load balancer 120 is a computing device including a processor and a memory. The memory stores computer code instructions and a data structure 121 associated with the anycast address to which the data packet 10 is originally addressed. The computer code instructions include a software load balancer for load balancing data packets associated with the anycast address. In some implementations, the software load balancer is a virtual machine. The load balancer 120 may include multiple software load balancers associated with the same anycast address and/or multiple software load balancers associated with multiple anycast addresses. The data structure 121 is employed by the load balancer 120 to map data packets associated with the anycast address to corresponding application instances 131 residing in the clusters 130. In some instances, each load balancer 120 includes a separate data structure for each anycast address served by the same load balancer 120. The processor is configured to execute the computer code instructions stored in the memory. A person of ordinary skill in the art should appreciate that each load balancer 120 may include multiple processors and/or multiple memories. The load balancers 120 include a computer server, an end-point server that terminates TCP connection, other electronic devices configured to perform load balancing as described herein, or combinations thereof.

[0020] In some instances, the load balancers 120 are distributed across different geographical areas of the system 100. For example the load balancers 120a, 120b, and 120c, serve corresponding clusters 130a, 130b, and 130c, respectively. In other instances, a load balancer 120 may serve more than one cluster 130. A load balancer 120 may serve a zone of the system 100. A zone herein refers to one or more clusters 130 of the system 100 which are related, for example, based on corresponding geographical locations or other criteria. For instance, a zone may include multiple clusters 130 located at close geographic proximities to each other such as within a same data center. In other instances, a zone may include more distant clusters 130 inter-connected through relatively high speed communication links. Also, a zone may be served by a pool of load balancers.

[0021] In instances where the received data packet 10 is redirected by the anycast node 110 to a pool of load balancers 120, equal-cost multi-path (ECMP) routing may be employed to forward the received data packet to a specific load balancer 120 of the pool of load balancers 120. For example, a device receiving the data packet 10 generates an integer hash value

based on header fields of the data packet 10, e.g., the packet's five-tuple, including protocol, source IP address, destination IP address, source port, and destination port. The receiving device then determines the destination load balancer 120 based on the generated hash value modulo a size of a table including identifications of the load balancers 120. In some implementations, the receiving device determines the destination load balancer 120 using the same or substantially similar algorithm as used by the load balancer 120 as described further below. The receiving device then delivers the data packet 10 to the destination load balancer 120 by either rewriting the packet's layer-two Ethernet header to point to the destination load balancer 120, or by encapsulating the data packet 10 with a generic routing encapsulation (GRE) header and an outer IP header. The same approach may be employed by a load balancer 120 including multiple software load balancers serving the packet's destination IP address to forward the received data packet 10 to one of the multiple software load balancers.

[0022] Upon receiving the data packet 10, the receiving load balancer 120, or a software load balancer thereon, generates a hash value based on one or more header fields of the data packet 10, e.g., the packet's five-tuple including protocol, source IP address, destination IP address, source port, and destination port. In some implementations, the data structure 121 includes entries associated with a group of application instances 131 served by the receiving load balancer 120. The group of application instances 131 correspond to an application or service associated with the anycast address to which the data packet 10 was addressed when arriving at the system 100. The receiving load balancer 120 then uses the generated hash value and the data structure 121 to determine a destination IP address of an application instance 131 of a group of application instances served by the receiving load balancer 120. The receiving load balancer 120 then forwards the data packet 10 to the determined application instance 131.

[0023] In some implementations, the data structure 121 is designed in a way that the frequency with which each application instance 131 is included in the data structure 121 reflects a weight associated with the same application instance 131. The weight associated with the application instance 131 is indicative of a load balancing metric such as capacity, processing time, or other criteria relevant in deciding to which application instance 131 the data packet is to be forwarded. That is, the weight associated with each application instance 131 may reflect how busy or how free the application instance 131 is. Alternatively, the weight associated with a given application instance 131 may reflect how fast or how slow processing the data packet 10 would be if the same application instance 131 is used. Capacity of each application may be defined in terms of packets per second, connections, or synchronize (SYN) packets per second, bandwidth, or the like. Processing time may be defined in terms of round trip time (RTT), or other metrics known in the art. In the data structure 121, the more frequent an application instance 131 is included, the more likely the same application instance 131 is to be selected for processing the request associated with the data packet 10.

[0024] The receiving load balancer 120 uses the generated hash value to select an entry of the data structure 121. For instance, the receiving load balancer 120 may select the data structure entry with index equal to the generated hash value modulo a size of the data

structure or the size of an element portion of the data structure. Each entry of the data structure 121 is indicative of a corresponding application instance 131. For instance, each entry of the data structure 121 includes a destination IP address of a corresponding application instance 131. The receiving load balancer 120 then forwards the data packet 10 to the application instance 131 with the destination IP address obtained from the selected data structure entry.

[0025] In some implementations, the data structure 121 is a single table with each row or alternatively each column, including entries for application instances associated with a corresponding location, e.g., a cluster 130 or zone, of the system 100. As such, the receiving load balancer 120 may first select a row, or alternatively a column, and then select an entry within the selected row, or alternatively the selected column, based on the generated hash value. In some implementations, the data structure 121 includes multiple tables with each table corresponding to a location, e.g., a zone or cluster 130, of the system 100. In such a case, the receiving load balancer 120 may first select a table, and then select an entry within the selected table based on the generated hash value.

[0026] The receiving load balancer 120 may select a table, or sub-table such as a row or a column, corresponding to a location of the system 100 in different ways. For instance the selection may be based on a map associating IP subnets to corresponding locations of the system 100. Such map may be defined by observing at what location of the system 100 traffic from each IP subnet usually arrives. Alternatively, the map may be defined based on assignments made by administrators of the system 100, or based on RTT and distances between IP subnets and locations 130 of the system 100. The receiving load balancer 120 looks up the source IP address associated with the received data packet 10 in the map, and gets back either a location or a weighted list of locations. If a weighted list of location is retrieved from the map, the receiving load balancer may select a location from the list based on the corresponding weights and another hash value generated using one or more header fields, e.g., the five tuple, of the data packet 10. However, if no weighted list is used, the receiving load balancer 120 checks whether the closest location indicated in the map has enough available capacity to handle the data packet 10. If yes, the closest location is selected, otherwise the next closest location as indicated in the map is checked and so on. Given that each location is associated with a corresponding table, or sub-table, the receiving load balancer 120 selects a table, or sub-table, for use in determining a destination application instance 131 for the data packet 10 when selecting a location from the map. A person skilled in the art should appreciate that the data structure 121 may alternatively include one or more trees or any other data structures known in the art instead of one or more tables.

[0027] The receiving load balancer 120 may also check a connection table, prior to selecting an application instance 131, to determine if the received data packet 10 corresponds to a flow or session associated with an existing connection. If the received data packet 10 is found to be associated with an existing connection, the data packet 10 is forwarded to an application instance associated with the connection in the connection table. Otherwise, the receiving load balancer 120 determines an application instance 131 from the data structure 121 and forwards

the data packet 10 to the determined application instance 131. Also, upon receiving the data packet 10, the load balancer 120 may further de-encapsulate the data packet 10 if the later has been previously encapsulated one or more times. As such, the receiving load balancer 120 de-encapsulates the data packet 10 until the inner packet is reached, and then retrieves any header fields for use in determining an application instance 131 to process the data packet 10.

[0028] Each cluster 130 includes one or more computer servers, e.g., content delivery servers and/or application servers, that maintain application instances associated with one or more services. In Figure 1, each of the clusters 130a - 130c, includes a number of application instances 131 associated with the application or service accessed through the anycast address to which the data packet 10 was addressed when arriving at the system 100. The application instances 131 in a given cluster 130 may reside in a single computer server, or may be distributed among multiple computer servers of the same cluster 130. The application instances 131 are addressed through corresponding destination IP addresses. An application instance as referred to herein includes a copy of a server application serving requests of end-users of a web page or an application, such as an email application, game application, social media application, calendar application, or any other online application. The application instance 131 receiving the data packet 10 processes the request associated with the data packet 10.

[0029] Figure 2 shows a flowchart describing an implementation of a process 200 of processing anycast data packets performed by the single-layer load balancing system 100. The process 200 includes the processes of receiving by an anycast node 110 a data packet 10 addressed to an anycast address (stage 210), forwarding by the anycast node 110 the received data packet 10 to a load balancer 120 (stage 220), and determining by the load balancer 120 if the data packet is associated with a previously served flow or session (stage 230). If the data packet is determined to be associated with a previously served flow or session, the process 200 includes forwarding the data packet 10 to application instance associated with the previously served flow or session (stage 240). Otherwise, the process 200 includes selecting a sub-data structure from multiple sub-data structures (stage 250), determining an application instance 131 based on the selected sub-data structure and one or more header fields of the data packet 10 (stage 260), and forwarding the data packet to the determined application instance (stage 270).

[0030] When end-users request or consume an online service associated with an anycast address, corresponding data packets sent from the end-users are addressed to the same anycast address. One or more anycast nodes 110 receive the data packets addressed to the anycast address (stage 210). For instance, each data packet addressed to the anycast address is received by the closest anycast node 110 to the source of the data packet. Upon receiving a data packet addressed to the anycast node (stage 210), a receiving anycast node may check a connection table to determine whether the received data packet is associated with a previously served flow or session, e.g., an already established connection. If the data packet is determined to be associated with an existing flow or session, the receiving anycast node 110 forwards the data packet to a next hop associated with the existing flow or session.

The checking of the connection table is optional as it may be carried out by another network element or device, other than the receiving anycast node 110 or it may be skipped entirely.

[0031] If the data packet is determined not to be associated with an existing flow or session such as a synchronize (SYN) data packet or a data packet where flow had previously been by a different anycast node, or no checking is performed by the receiving anycast node 110, the anycast node forwards the data packet to a load balancer (LB) 120 (stage 220). In forwarding the data packet (stage 220), the receiving anycast node 110 may determine the load balancer 120 based on one or more header fields, e.g., source IP address, destination IP address, and/or source port, of the data packet and a sub-data structure associated with the receiving anycast node 110. Also, the anycast node 110 may de-capsulate the data packet and/or encapsulate it with one or more new headers before forwarding the data packet to the load balancer 120.

[0032] Upon receiving the data packet, the load balancer 120 may check a connection table to determine whether the received data packet is associated with a previously served flow or session, e.g., an already established connection (stage 230). If the data packet is determined to be corresponding to an existing flow or session, the LB 120 forwards the data packet to the application instance serving the existing flow or session (stage 240). The checking of the connection table is optional as it may be carried out by another network element or device, other than the LB 120 or it may be skipped entirely.

[0033] If the data packet is determined not to be associated with an existing flow or session such as a synchronize (SYN) data packet or a data packet where flow had previously been by a different anycast node, or no checking is performed by the LB 120, the LB selects a sub-data structure from a data structure 121 maintained by the LB 120 (stage 250). For instance, if the data structure 121 is a single table, the selected sub-data structure may be a row or column of the table. In other instances where the data structure 121 includes multiple tables, the selected sub-data structure is table of the multiple tables. The selection of the sub-data structure may be based on header field(s) of the data packet and a sub-data structure between IP subnets and corresponding locations of the system 100. Alternatively, the selection may be based on header field(s) of the data packet and a list of the sub-data structures reflecting a weight for each sub-data structure in the list. Each sub-data structure of the data structure 121 represents a redundant list of application instances 131 associated with a corresponding location, e.g., a cluster 130 or zone, of the system 100. The frequency with which an application instance 131 is included in a corresponding sub-data structure is dependent on a weight value associated with the same application instance 131.

[0034] The LB 120 then determines an application instance 131 from the selected sub-data structure using one or more header fields of the data packet (stage 260). For instance, the LB 120 generates a hash value using the one or more header fields of the data packet and then uses the generated hash value to determine an entry of the sub-data structure. The LB 120 may calculate the generated hash value modulo the size of the sub-data structure and use the result as an index for the entry selected. A person of ordinary skill in the art should appreciate

that the header field(s) of the data packet may be used in different ways to identify an entry of the sub-data structure for selection. Each entry of the selected sub-data structure includes an IP address, e.g., a destination IP address, associated with a corresponding application instance 131. The LB 120 then forwards the data packet to the determined application instance 131 (stage 270), where the request associated with the data packet is served.

[0035] Figure 3 shows a block diagram of an implementation of a two-layer load balancing system 300 for load balancing data traffic addressed to an anycast address. The system 300 includes anycast redirector nodes 310a - 310c, also referred to hereinafter either individually or collectively as anycast node(s) 310, first-layer load balancer devices 320a - 320c, also referred to hereinafter either individually or collectively as first-layer load balancer(s) 320, second-layer load balancers 325a - 325c, also referred to hereinafter either individually or collectively as second-layer load balancer(s) 325, and multiple server clusters, e.g., clusters 330a - 330c also referred to hereinafter either individually or collectively as cluster(s) 330. Each of the clusters 330a - 330c, includes a number of application instances 331 associated with a service accessed through the anycast address.

[0036] An anycast node 310 is configured to forward a received data packet 10, addressed to a corresponding anycast address, to a first-layer LB 320 of one or more first-layer LBs 320. The anycast node 310 determines the first-layer LB 320 for forwarding the data packet based on one or more header fields, e.g., source IP address, destination IP address, and/or source port, of the data packet 10 and a sub-data structure associated with the anycast node 310. Also, the anycast node 310 may de-capsulate the data packet and/or encapsulate it with one or more new headers before forwarding the data packet to the selected first-layer LB 320. Furthermore, the anycast node 310 may check a connection table upon receiving the data packet similar to the anycast node 110 described in relation to Figure 1.

[0037] In some implementations, each first-layer LB 320 includes a first data structure 321 for mapping received data packets to respective second-layer LBs 325. In some instances, the first data structure 321 includes a redundant list of the second-layer LBs 325 such that the frequency with which each second-layer LB 325 is included in the first data structure depends on a weight value associated with the same second-layer LB 325 or a corresponding location. For instance, the weights may reflect the available capacity at each corresponding location, the RTT to each corresponding location, another load balancing criterion, or combinations thereof. The first-layer LB 320 selects a second-layer LB 325 from the corresponding first data structure 321 based on one or more header fields of the data packet 10. The first-layer LB 325 then forwards the data packet 10 to the selected second-layer LB 325. In some instances, the first-layer LB 320 generates a hash value using the header field(s) of the data packet 10 and selects the second-layer LB 325 from the first data structure 321 based on the generated hash value. In other instances, the selection of the second-layer LB 325 by the first-layer LB 320 may be performed similar to the selection of a sub-data structure (stage 250) as described in relation with Figure 2. The selection of a second layer LB 325 by the first-layer LB 320 may be dependent on checking a connection table by the first-layer LB 320 (similar to stage 230, stage 240 and stage 250 in Figure 2).

[0038] Each second-layer LB 325 is associated with a location, e.g., zone or cluster 330, of the system 300, and includes a corresponding second data structure 326 associated with the same location of the system 300. The second data structure 326 includes a redundant list of application instances associated with the same location. The frequency with which each application instance is included in the second data structure 326 reflects a weight value associated with the same application instance. Such weights may reflect the available capacity at each application instance, the RTT to each application instance, other load balancing criteria, or combinations thereof. The selected second-layer LB 325 receives the data packet 10 and determines an application instance from the corresponding second data structure 326 based on one or more header fields of the data packet 10. For instance, the first-layer LB 320 generates a hash value using the header field(s) of the data packet 10 and selects the second-layer LB 325 from the first data structure 321 based on the generated hash value. In some implementations, each entry of the second data structure 326 includes an IP address, e.g., a destination IP address, associated with a corresponding application instance. The generated hash value can be used as an index of an entry in the second data structure in selecting an IP address. The second-layer LB 325 then forwards the data packet 10 to the determined application instance.

[0039] Each cluster 330 includes one or more computer servers, e.g., content delivery servers and/or application servers, that maintain application instances associated with one or more services. In Figure 3, each of the clusters 330a - 330c, includes a number of application instances 331 associated with the application or service accessed through the anycast address to which the data packet 10 was addressed when arriving at the system 300. The application instances 331 in a given cluster 330 may reside in a single computer server, or may be distributed among multiple computer servers of the same cluster 330. The application instances 331 are addressed through corresponding destination IP addresses. The application instance 331 receiving the data packet 10 processes the request associated with the data packet 10. As illustrated in Figure 3, each cluster 330, or zone, of the system 300 may be served by one or more second-layer LBs 325. For instance, while the cluster 330a is served by the second-layer LB 325a, the cluster 330b is served by two second-layer LBs 325b and 325b'.

[0040] Figure 4 shows a block diagram of another implementation of the two-layer load balancing system 300. For convenience of illustration, the block diagram shown in Figure 4 depicts only a single anycast node 310, two first-layer LBs 320a and 320b, two second-layer LBs 325a and 325b, and two clusters 330a and 330b, of the system 300. The system 300 also includes a first controller 312, a second controller 322, a third controller 327, and a global load balancing controller 350.

[0041] The first controller 312 includes a network element, a computer server, or another electronic device. The first controller 312 configures the anycast node 310, with information such as the source IP map, connection table, and/or other information employed by the anycast node 310 in handling received data packets 10 addressed to a corresponding anycast address. The first controller 312 may acquire such information from one or more data bases or

other devices of the system 300, and provide corresponding updates to the anycast node 310.

[0042] The second controller 322 includes a network element, a computer server, or another electronic device. The second controller 322 configures first-layer LBs 320a and 320b, by providing information such as the weights associated with each location, e.g., cluster 330 or zone, of the system 300, routing information for routing data packets to locations of the system 300, a connection table, and/or other information employed by the first-layer LB 320 in handling received data packets 10. The second controller device 322 may acquire such information from the global load balancing controller 350, one or more databases, or other devices of the system 300, and provide corresponding updates to the first-layer LB 320.

[0043] The third controller 327 includes a network element, a computer server, or another electronic device. The third controller 327 configures the second-layer LBs 325a and 325b, by providing information such as the weights associated with each application instance 331 in a corresponding location of the system 300, routing information for routing data packets 10 to application instances 331, and/or other information employed by the second -layer LB 320 in handling received data packets 10. The third controller 327 may acquire such information from the one or more databases or other devices of the system 300, and provide corresponding updates to the second-layer LB 320.

[0044] In Figure 4, the continuous lines between different components of the system 300 indicate the data packet path (also referred to as the data plane), the dashed lines indicate the flow of configuration information/instructions (also referred to as the control plane), and the break lines indicate feedback paths (which may also be part of the control plane). In some implementations, each time the first-layer LB 320 forwards a data packet to a location, or a corresponding second-layer LB 325, of the system 300, the first-layer LB 325 reports the forwarding to the global load balancing controller 350. Also, the second-layer LB 325 may report the forwarding of the data packet 10 to the global load balancing controller 350 and/or to another device of the system 300. The global load balancing controller 350 uses the information reported by the first-layer LB 320 to update weights associated with different locations of the system 300. The information reported by the second-layer LB 325 is used by the global load balancing controller 350, or a device local at the location associated with the second-layer LB 325, to update weights associated with different application instances 331. In some implementations, the clusters 330 can be configured to provide status information of its servers or application instances 331 to the global load balancing controller 350 and/or other devices in the system 300. The global load balancing controller 350 or another device in the system 300 can use the information reported by the clusters 330 to update weights associated with different application instances 331. In some implementations, the global load balancing controller 350 also obtains data from routers and/or network elements regarding link congestions.

[0045] Figure 5 shows a flowchart describing an implementation of a process 500 of processing anycast data packets performed by the two-layer load balancing system 300. The process 200 includes the processes of receiving by an anycast node 310 a data packet 10

addressed to an anycast address (stage 510), forwarding by the anycast node 310 the received data packet 10 to a first-layer LB 320 (stage 520), determining by the first-layer LB 320 if the data packet is associated with a previously served flow or session (stage 530), if the data packet 10 is determined to be associated with a previously served flow or session forwarding the data packet 10 to an application instance associated with the previously served data flow or session (stage 540), otherwise selecting a second-layer LB 325 based on a first data structure maintained by the first-layer LB 320 and one or more header fields of the data packet 10 (stage 550), forwarding the data packet to the selected second-layer LB 325 (stage 560), determining by the selected second-layer LB 325 an application instance 331 based on a second data structure maintained by the second-layer LB 325 and one or more header fields of the data packet 10 (stage 570), and forwarding the data packet to the determined application instance (stage 580).

[0046] The stages 510 - 540 of the process 500 are similar to the stages 210 - 240 of the process 200 described in relation with Figure 2, except that the first-layer LB 325 (shown in Figures 3 and 4) is used instead of the load balancer 120 (shown in Figure 1). Also, stages 530 and 540 are optional since the checking of the connection table may be performed by another device, other than the first-layer LB 320. The first layer LB 320 receiving the data packet 10 selects a second layer LB 325 based on a first data structure 321 maintained by the first-layer LB 320 and one or more header fields of the data packet 10 (stage 550). The first data structure 321 includes a redundant list of second-layer LBs 325 or corresponding locations, e.g., clusters 330 or zones, in the system 300. The frequency with which each second-layer LB 325, or a corresponding location, is included in the first data structure 321 depends on weight value associated with same corresponding location. The weights may reflect the available capacity at each corresponding location, the RTT to each corresponding location, another load balancing criterion, or combinations thereof. In some instances, the first-layer LB 320 generates a hash value using the header field(s) of the data packet 10 and selects the second-layer LB 325 from the first data structure 321 based on the generated hash value. For instance, the first-layer LB 320 uses the generated hash value modulo the size of the first data structure as an index of an entry to be selected from the first data structure. The first-layer LB 320 then forwards the data packet to the selected second-layer LB 325, or to a second-layer LB 325 associated with the selected location (stage 560).

[0047] The second-layer LB 325 receiving the data packet 10 determines an application instance 331 based on a second data structure 326 maintained by the second-layer LB 325 and one or more header fields of the data packet 10 (stage 570). The second data structure 326 includes a redundant list of application instances 331. The frequency with which each application instance 331 is included in the second data structure 326 depends on weight values associated with the application instances 331. The weights may reflect the available capacity at each corresponding application instance 331, the RTT to each corresponding application instance 331, other load balancing criteria, or combinations thereof. The second-layer LB 325 generates a hash value using the header field(s) of the data packet 10 and determines the application instance 331 from the second data structure 326 based on the generated hash value. For instance, the second-layer LB 325 uses the generated hash value

modulo the size of the second data structure 326 as an index of an entry to be selected from the second data structure 326. The second-layer LB 325 then forwards the data packet 10 to the determined application instance 331 (stage 580).

[0048] Figure 6 shows a flowchart describing an implementation of a process 600 for generating a redundant list employed by load balancers. For instance, given a list of entities each being associated with a corresponding weight value, a prime number, larger than the number of entities in the list, is chosen as the size of the redundant list to be generated. A processor executing the process 600 selects an offset value and a step value for each entity in the given list of entities (stage 610). The processor then selects an entity from the given list of entities (stage 620). The selection can be done based on a pseudorandom permutation of identifications (such as names, portions of names, identification strings, or the like) of the entities. The processor compares the number of entries already included in the redundant list for the selected entity to a corresponding dynamic threshold value (stage 630). For instance, the dynamic threshold value is defined as the weight corresponding to the selected entity multiplied by an iteration number. If the number of entries already included is found to be smaller than the dynamic threshold value at (stage 630), the processor checks if the position indicated by the offset value in the redundant list to be generated is empty (stage 640). If the position indicated by the offset value is not found empty, the processor updates the offset value by incrementing it with the corresponding step value and truncating the incremented value modulo the size of the redundant list (stage 650). The processor then checks if the position indicated by the updated offset value is empty (stage 640). The stages 650 and 640 are repeated until an empty position is detected. If at any point, the result of the process (stage 640) indicates an empty position, the processor adds an entry corresponding to the selected entity at the empty position (stage 660). Once an entry for the selected entity is added in the redundant list (stage 660), or the number of already added entries for the selected entity is found to be larger than the dynamic threshold at (stage 630), the processor then checks if all the entities in the given list have been processed in the current iteration (stage 670). If not all entities in the given list have been processed, the processor selects a new entity from the given list for processing (stage 620). otherwise, the processor increments the iteration number (stage 680), and then selects an entity from the given list for processing (stage 620). The processor iterates through the stages depicted in Figure 6 until the redundant list is full.

[0049] The process 600 generates the redundant list as one-dimensional table, e.g., a row or column, a tree, a sub-tree, or any other data structure. The entities in the given list may be application instances 131 or 331, load balancers 120 or 325, or locations of a load balancing system 100 or 300. In the generated redundant list, the frequency with which each entity is repeated is dependent on the weight value corresponding to the same entity. In instances where the size of the redundant list is chosen to be much larger than the total number of entities in the give list, e.g., 100 times larger, the frequencies with which the entities are included in the redundant list becomes almost proportional to the corresponding weights. The larger the sized of the generated redundant list relative to the number of entities in the given list, the more accurate is the proportionality between the frequencies and the weights of the same entities.

[0050] Figure 7 shows illustrations of data structures employed by systems in Figures 1 and 3. Considering three application instances A_1, A_2, and A_3 of a cluster A, e.g., clusters 130a or 330a, and three application instances B_1, B_2, and B_3 of a cluster B, e.g., clusters 130b or 330ba, each row of the table in Figure 7 depicts a sample output, with size equal to 17, of the process 600 (shown in Figure 6) generated based on corresponding sets of weights. For instance, the first row corresponds to an output where the weights associated with each application instance A_1, A_2, and A_3 are equal to 1.0 and the weights associated with the application instances B_1, B_2, and B_3 are equal to 0.0. Going from any row to the next, the weight value for the application instances A_1, A_2, and A_3 is decremented by 0.1, and the weight value for the application instances B_1, B_2, and B_3 is incremented by 0.1.

[0051] As the weights associated with the application instances change slowly from one row to the next in the table shown in Figure 7, only few changes are observed from one row to the next. In fact, by examining the columns of the table of Figure 7, one can see that only few changes occur across each column as the weights slightly change from one row to the next.

[0052] Consider a scenario where the cluster 130a of Figure 1 includes the application instances A_1, A_2, and A_3 and that each of these application instances is experiencing an overflow. The overflow demand, e.g., 20%, is to be redirected to the next closest cluster with available capacity, e.g., cluster 130b of Figure 1 including application instances B_1, B_2, and B_3. As such, a weight value of 0.8 is associated with each of the application instances A_1, A_2, and A_3, and a weight value of 0.2 is associated with each of the application instances B_1, B_2, and B_3 with respect to anycast traffic associated with the cluster 130a of Figure 1. In such, a case the third row of the table in Figure 7 represents a sample of a sub-data structure associated with the cluster 130a of Figure 1. Also, given that the application instances B_1, B_2, and B_3 have enough bandwidth to serve anycast traffic coming to cluster 130b, these application instances will all have equal weight of 1.0 when it comes to anycast traffic associated with the cluster 130b shown in Figure 1. As such, the last row of the table in Figure 7 represents a sample of the sub-data structure associated with the cluster 130b of Figure 1.

[0053] Comparing any pair of consecutive rows, one can observe that only few entries change from one row to the next. Such observation indicates that a small change in the weight values (the weight values change by ± 0.1 from one row to next one) has a slight effect on data paths. Adding or removing one application instance associated with a small weight value would also have slight effect on the entries of the respective row or sub-data structure. The way rows or sub-data structures are generated (using the process 600 shown in Figure 6) results in consistency in assigning new requests to respective application instances 131 or 331 (shown in Figures 1 and 3). In some implementations, small changes in weight values and/or adding/removing application instances associated with small weights would not result in rerouting of data traffic (or re-computing routing tables at network routers or other network elements) since such changes have slight effect on load balancing data traffic.

[0054] Considering a similar scenario, e.g., with the same weights as discussed in the previous

paragraph, where the application instances are associated with the cluster 330a of Figure 3, and the application instances B_1, B_2, and B_3 are associated with the cluster 330b of Figure 3, then the third row of the table in Figure 7 represents a sample of the second data structure 326 associated with the cluster 330a of Figure 3. Also, the last row of the table in Figure 7 represents a sample of the second data structure 326 associated with the cluster 330b of Figure 3.

[0055] Implementations of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software embodied on a tangible medium, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the subject matter described in this specification can be implemented as one or more computer programs embodied on a tangible medium, i.e., one or more modules of computer program instructions, encoded on one or more computer storage media for execution by, or to control the operation of, a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-readable storage substrate, a random or serial access memory array or device, or a combination of one or more of them. The computer storage medium can also be, or be included in, one or more separate components or media (e.g., multiple CDs, disks, or other storage devices). The computer storage medium may be tangible and non-transitory.

[0056] The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0057] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular implementations of particular inventions. Certain features that are described in this specification in the context of separate implementations can also be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation can also be implemented in multiple implementations separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a sub-combination or variation of a sub-combination.

[0058] References to "or" may be construed as inclusive so that any terms described using "or" may indicate any of a single, more than one, and all of the described terms. The labels "first," "second," "third," and so forth are not necessarily meant to indicate an ordering and are generally used merely to distinguish between like or similar items or elements.

[0059] Thus, particular implementations of the subject matter have been described. Other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results. In addition, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking or parallel processing may be utilized.

REFERENCES CITED IN THE DESCRIPTION

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

Patent documents cited in the description

- [US2010302940A \[0003\]](#)
- [US7355977B \[0004\]](#)
- [US2011145390A \[0005\]](#)

P A T E N T K R A V

1. System til belastningsbalancing af anycast-trafik i et kommunikationsnetværk omfattende:

et første sæt af belastningsbalanceringsanordninger (LB-anordninger), hvor hver LB-anordning af det første sæt af LB-anordninger er konfigureret til:

at opretholde en første datastruktur, som indbefatter indlæsninger, der er tilknyttet applikationsinstanser i en gruppe af applikationsinstanser, som er betjent af LB-anordningen af det første sæt af LB-anordninger, hvor hyppigheden, med hvilken hver betjent applikationsinstans er indbefattet i den første datastruktur, er indikerende for en vægtværdi, der er tilknyttet en kapacitet af den tilsvarende betjente applikationsinstans;

ved modtagelse af en datapakke, der er adresseret en anycast-adresse, ved en LB-anordning af det første sæt af LB-anordninger, at generere en første hashværdi baseret på ét eller flere første startfelter af den modtagne datapakke,

at identificere en virtuel internetprotokoladresse (IP-adresse) af én af de betjente applikationsinstanser baseret på den første genererede hashværdi ved anvendelse af den første datastruktur; og

at videresende datapakken til den identificerede applikationsinstans;

et andet sæt af belastningsbalanceringsanordninger (LB-anordninger), hvor hver LB-anordning af det andet sæt af LB-anordninger er konfigureret til:

20 at opretholde en anden datastruktur, som indbefatter indlæsninger, der er tilknyttet de respektive LB-anordninger af det første sæt, hvor hyppigheden, med hvilken hver LB-anordning i det første sæt af LB-anordninger er indbefattet i den anden datastruktur, er indikerende for en vægtværdi, der er tilknyttet den tilsvarende LB-anordning af det første sæt;

25 at generere en anden hashværdi ved modtagelse af datapakken baseret på ét eller flere anden startfelter af den modtagne datapakke;

at identificere en LB-anordning af det første sæt af LB-anordninger ved anvendelse af den anden datastruktur baseret på den genererede anden hashværdi; og

30 at videresende datapakken til den identificerede LB-anordning af det første sæt; og en flerhed af anycast-knudepunkter, som er tilknyttet anycast-adressen, og som er konfigureret til:

35 at slå pakkens kildeadresse op i et kildeadressekort, som deles af flerheden af anycast-knudepunkter, ved modtagelse af datapakken for at bestemme en placering, som er tilknyttet kildeadressen, hvor pladsen er indikerende for en LB-anordning af det andet sæt af LB-anordninger, og at videresende den modtagne datapakke til LB-anordningen.

2. System ifølge krav 1, hvor LB-anordningen af det andet sæt af LB-anordninger yderligere er konfigureret til at rapportere trafikbelastning, som er videresendt til hver LB-anordning af det første sæt af LB-anordninger, til et globalt LB-element.

3. Systemet ifølge krav 2, hvor det globale LB-element er konfigureret til at generere vægte, som er tilknyttet LB-anordninger af det første sæt af LB-anordninger i det mindste delvist baseret på modtagne rapporter af videresendte trafikbelastninger fra LB-anordninger af det andet sæt af LB-anordninger til LB-anordninger af det første sæt af LB-anordninger.

5 4. System ifølge et hvilket som helst af kravene 1 til 3, hvor vægtværdien, som er tilknyttet den tilsvarende LB-anordning af det første sæt, er indikerende for en kapacitet af gruppen af processeringsanordninger betjent af den tilsvarende LB-anordning af det første sæt, hvor applikationsinstanserne fortrinsvis er tilknyttet mindst én af en applikationsserver, en indholdsserver, og en virtuel maskine.

10 5. System ifølge et hvilket som helst af kravene 1 til 4, hvor hver LB-anordning af det andet sæt af LB-anordninger er yderligere konfigureret til:

at modtage en vægtværdi for én eller flere LB-anordninger i det første sæt af LB-anordninger; og

15 at generere den anden datastruktur baseret på de modtagne vægtværdier for den ene eller de flere LB-anordninger i det første sæt af LB-anordninger.

6. System ifølge krav 5, hvor hver LB-anordning af det andet sæt af LB-anordninger ved generering af den anden datastruktur er yderligere konfigureret til:

20 at vælge et antal tomme datastrukturplaceringer for hver LB-anordning af det første sæt af LB-anordninger, hvor antallet af valgte datastrukturplaceringer bestemmes baseret på vægtværdien, som svarer til LB-anordningen af det første sæt af LB-anordninger; og at indsætte en IP-adresse af LB-anordningen af det første sæt af LB-anordninger i hver af de valgte tomme datastrukturplaceringer,

25 hvor, fortrinsvis, hver LB-anordning af det andet sæt af LB-anordninger er yderligere konfigureret til at vælge antallet af tomme datastrukturplaceringer for hver LB-anordning af det første sæt af LB-anordninger baseret på en offsetværdi tilknyttet LB-anordningen af det første sæt af LB-anordninger.

7. System ifølge et hvilket som helst af kravene 1 til 6, hvor hver LB-anordning af det første sæt af LB-anordninger er yderligere konfigureret til:

30 at modtage en vægtværdi for hver applikationsinstans i gruppen af applikationsinstanser, der er betjent af LB-anordningen af det første sæt af LB-anordninger; og

at generere den første datastruktur baseret på de modtagne vægtværdier for applikationsinstanser i gruppen, som er betjent af LB-anordninger af det første sæt af LB-anordninger, hvor hver LB-anordning af det første sæt af LB-anordninger, fortrinsvis ved 35 genereringen af den første datastruktur, er yderligere konfigureret til:

at vælge et antal tomme datastrukturplaceringer for hver applikationsinstans af gruppen betjent af LB-anordningen af det første sæt af LB-anordninger, hvor antallet af valgte datastrukturplaceringer er bestemt baseret på vægtværdien svarende til applikationsinstansen; og

at indsætte en virtuel IP-adresse af applikationsinstansen i hver af de valgte tomme datastrukturplaceringer.

8. System ifølge krav 7, hvor hver LB-anordning af det første sæt af LB-anordninger er yderligere konfigureret til at vælge antallet af tomme datastrukturplaceringer for hver applikationsinstans af gruppen betjent af LB-anordningen af det første sæt af LB-anordninger baseret på en offsetværdi tilknyttet applikationsinstansen.

9. Fremgangsmåde til datatrafikbelastningsbalancing hvor:

hver belastningsbalanceringsanordning (LB-anordning), af et første sæt af LB-anordninger udfører de følgende trin:

10 at opretholde en første datastruktur indbefattende indlæsninger, der er tilknyttet applikationsinstanser i en gruppe af applikationsinstanser betjent af LB-anordningen af det første sæt af LB-anordninger, hvor hyppigheden, med hvilken hver betjent applikationsinstans er indbefattet i den første datastruktur, er indikerende for en vægtværdi, der er tilknyttet en kapacitet af den tilsvarende betjente applikationsinstans;

15 ved modtagelse, ved en LB-anordning af det første sæt af LB-anordninger, af en datapakke adresseret en anycast-adresse, at generere en første hashværdi baseret på ét eller flere første startfelter af den modtagne datapakke;

20 at identificere en virtuel internetprotokoladresse (IP-adresse)- af én af de betjent applikationsinstanser baseret på den genererede første hashværdi ved at anvende af den første datastruktur; og

25 at videresende datapakken til den identificerede applikationsinstans;

hvor hver LB-anordning af et andet sæt af LB-anordninger udfører de følgende trin:

at opretholde en anden datastruktur indbefattende indlæsninger tilknyttet respektive LB-anordninger i det første sæt, hvor hyppigheden, med hvilken hver LB-anordning i det første sæt af LB-anordninger er indbefattet i den anden datastruktur, er indikerende for en vægtværdi, der er tilknyttet den tilsvarende LB-anordning af det første sæt;

30 at generere en anden hashværdi ved modtagelse af datapakken baseret på ét eller flere anden startfelter af den modtagne datapakke, hvor datapakken modtages ved LB-systemet, som er adresseret anycast-adressen;

at identificere en LB-anordning af det første sæt af LB-anordninger, baseret på den genererede anden hashværdi ved at anvende den anden datastruktur;

at videresende datapakken til den identificerede LB-anordning af det første sæt af LB-anordninger; hvor en flerhed af anycast-knudepunkter tilknyttet anycast-adressen udfører følgende trin:

35 at slå pakkens kildeadresse op i et kildeadressekort, som deles af flerheden af anycast-knudepunkter, for at bestemme en placering tilknyttet kildeadressen, hvor placeringen er indikerende for en LB-anordning af det andet sæt af LB-anordninger, og at videresende den modtagne datapakke til LB-anordningen.

10. Fremgangsmåde ifølge krav 9, der yderligere omfatter at rapportere trafikbelastning, som er videresendt til hver LB-anordning af det første sæt af LB-anordninger, til et globalt LB-element ved hjælp af LB-anordningen af det andet sæt af LB-anordninger; som, fortrinsvis, yderligere omfatter, at generere vægte, der er tilknyttet LB-anordningerne af 5 det andet sæt af LB-anordninger, ved hjælp af det globale LB-element, i det mindste delvist baseret på modtagne rapporter af videresendte trafikbelastninger fra LB-anordninger af det andet sæt af LB-anordninger til LB-anordninger af det første sæt af LB-anordninger.

11. Fremgangsmåde ifølge et hvilket som helst af kravene 9 til 10, hvor vægtværdien, som er tilknyttet den tilsvarende LB-anordning af det første sæt af LB-anordninger, er 10 indikerende for en kapacitet af gruppen af applikationsinstanser betjent af den tilsvarende LB-anordning af det første sæt af LB-anordninger, og hvor applikationsinstanserne fortrinsvis er tilknyttet mindst én af en applikationsserver, en indholdsserver, og en virtuel maskine.

12. Fremgangsmåde ifølge et hvilket som helst af kravene 9 til 11, som yderligere 15 omfatter:

at modtage en vægtværdi for én eller flere LB-anordninger i det første sæt af LB-anordninger, ved hjælp af LB-anordningen af det andet sæt af LB-anordninger; og

at generere den anden datastruktur baseret på de modtagne vægtværdier for én eller flere LB-anordninger i det første sæt af LB-anordninger.

20 13. Fremgangsmåde ifølge krav 12, hvor trinnet at generere den anden tabel indbefatter:

at vælge et antal af tomme datastrukturplaceringer for hver LB-anordning af det første sæt af LB-anordninger, hvor antallet af valgte datastrukturplaceringer er bestemt baseret på vægtværdien tilsvarende til LB-anordningen af det første sæt af LB-anordninger; og 25 at indsætte en IP-adresse af LB-anordningen af det første sæt i hver af de valgte tomme datastrukturplaceringer,

30 hvor, fortrinsvis, at vælge antallet af tomme datastrukturplaceringer, for hver LB-anordning af det første sæt af LB-anordninger, indbefatter at vælge antallet af tomme datastrukturplaceringer baseret på en offsetværdi tilknyttet LB-anordningen af det første sæt af LB-anordninger.

14. Fremgangsmåde ifølge et hvilket som helst af kravene 9 til 13, som yderligere 35 omfatter:

at modtage en vægtværdi for hver applikationsinstans i gruppen af applikationsinstanser betjent af LB-anordningen af det første sæt af LB-anordninger ved LB-anordningen af det første sæt, og

at generere den første datastruktur baseret på de modtagne vægtværdier.

15. Fremgangsmåden ifølge krav 14, hvor trinnet at generere den første datastruktur indbefatter:

at vælge et antal af tomme datastrukturplaceringer, for hver applikationsinstans af gruppen betjent af LB-anordningen af det første sæt af LB-anordninger, hvor antallet af valgte datastrukturplaceringer bestemmes baseret på vægtværdien svarende til applikationsinstansen; og

5 at indsætte en virtuel IP-adresse af applikationsinstansen i hver af de valgte tomme datastrukturplaceringer,
hvor, fortrinsvis, at vælge antallet af tomme datastrukturplaceringer for hver applikationsinstans indbefatter at vælge antallet af tomme datastrukturplaceringer baseret på en offsetværdi, som er tilknyttet applikationsinstansen.

DRAWINGS

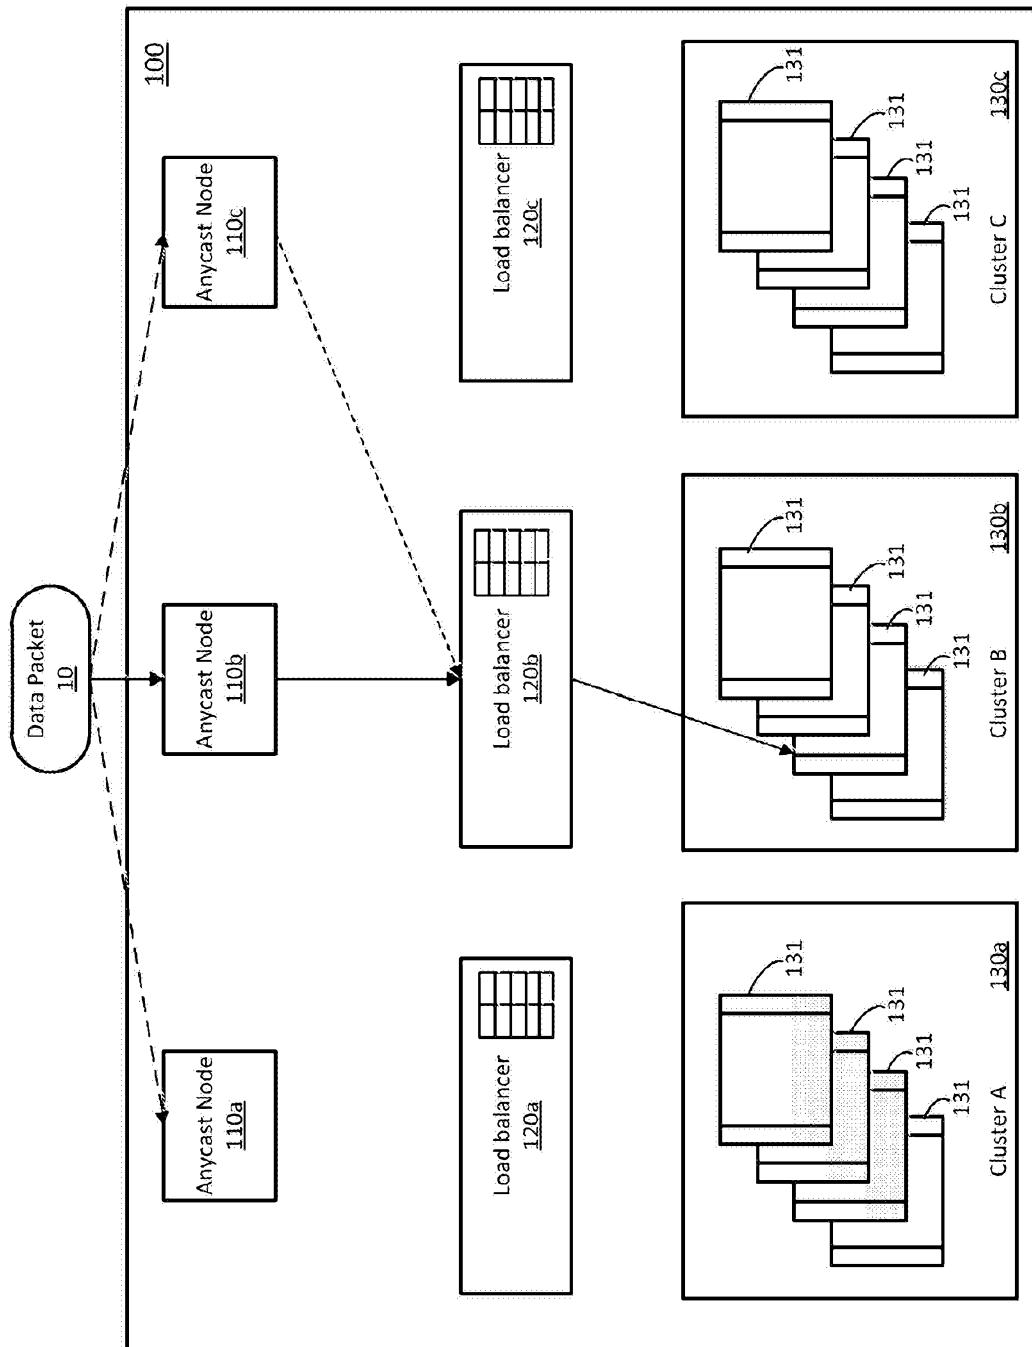


Figure 1

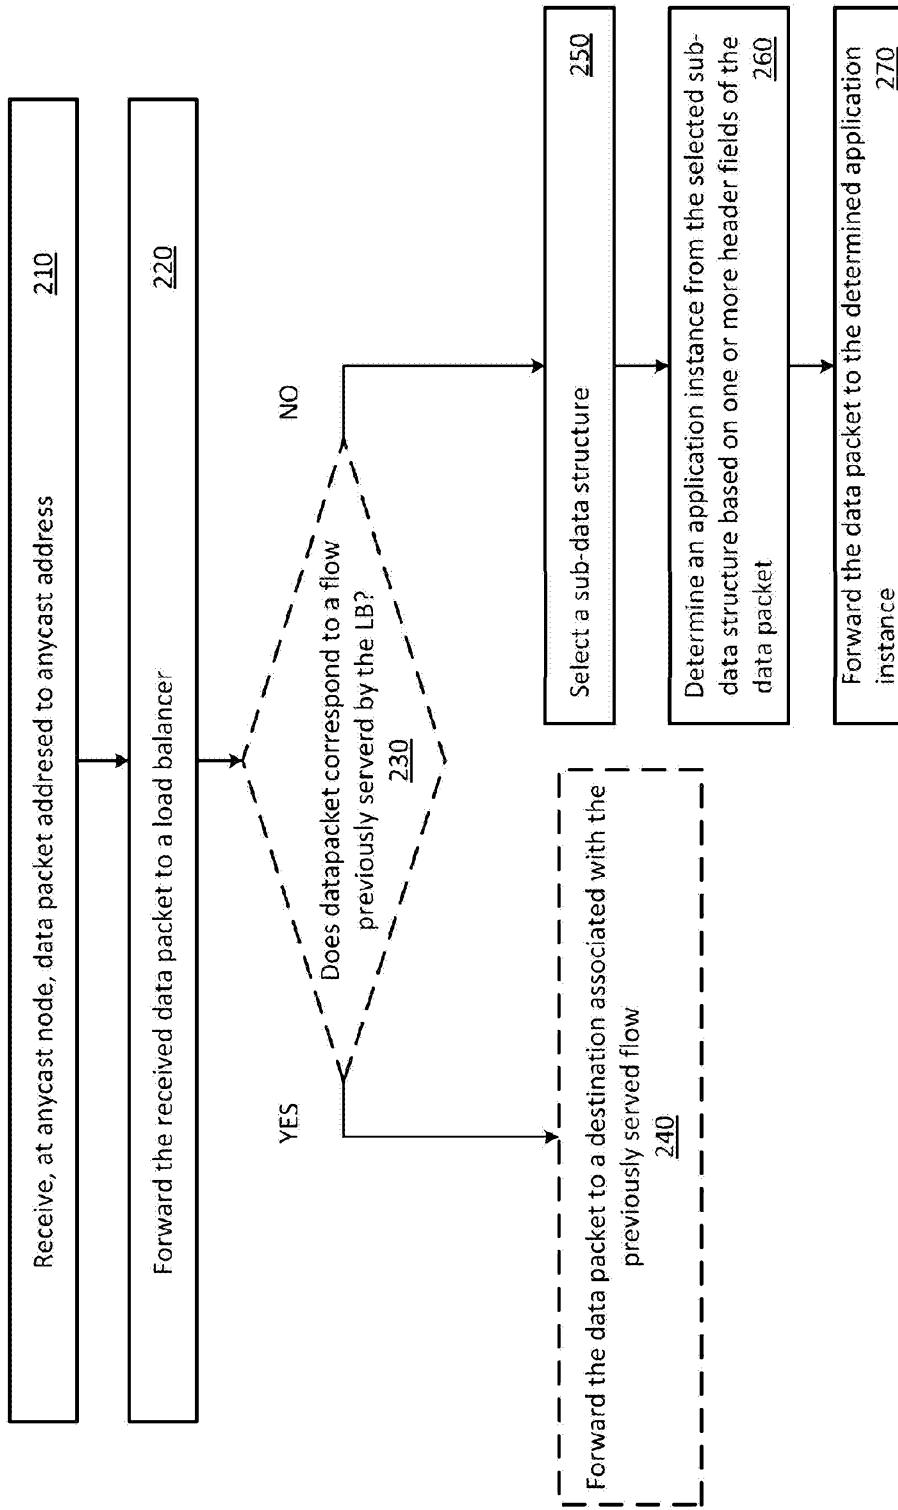


Figure 2

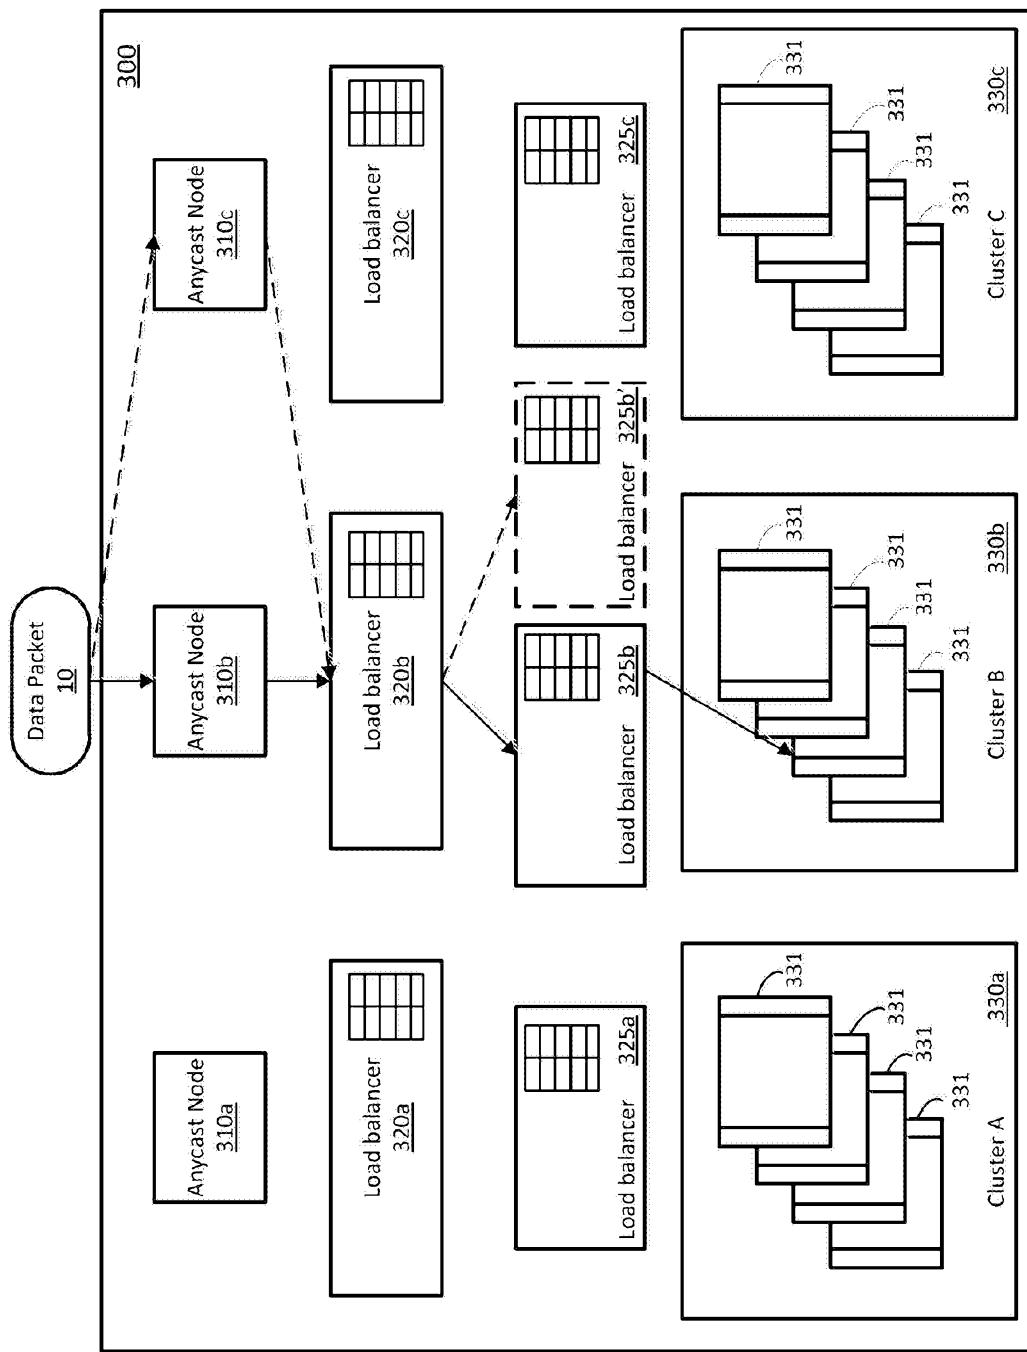


Figure 3

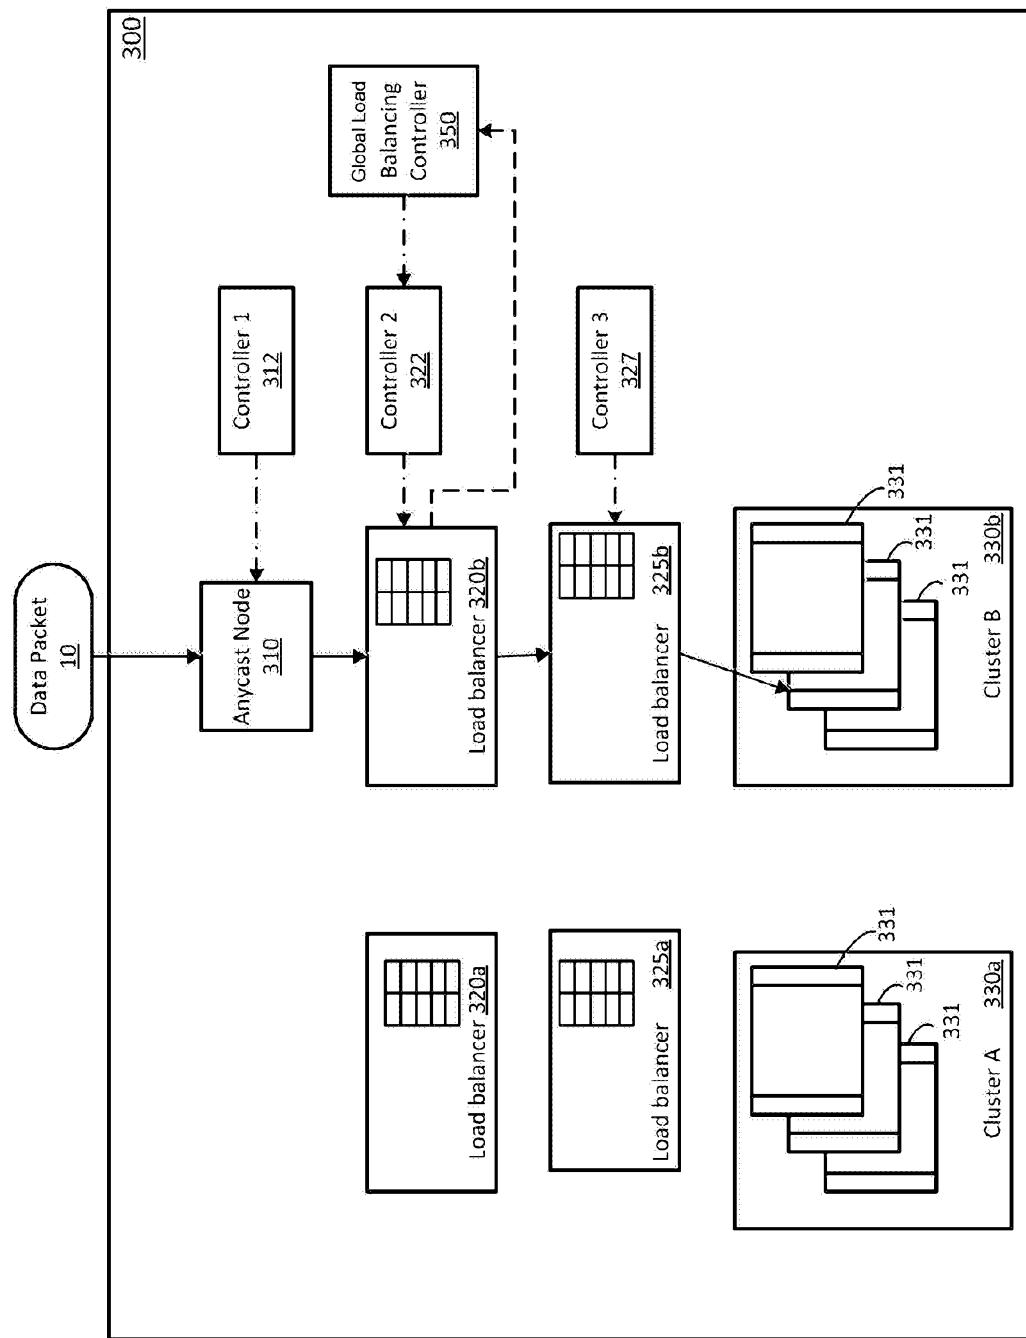


Figure 4

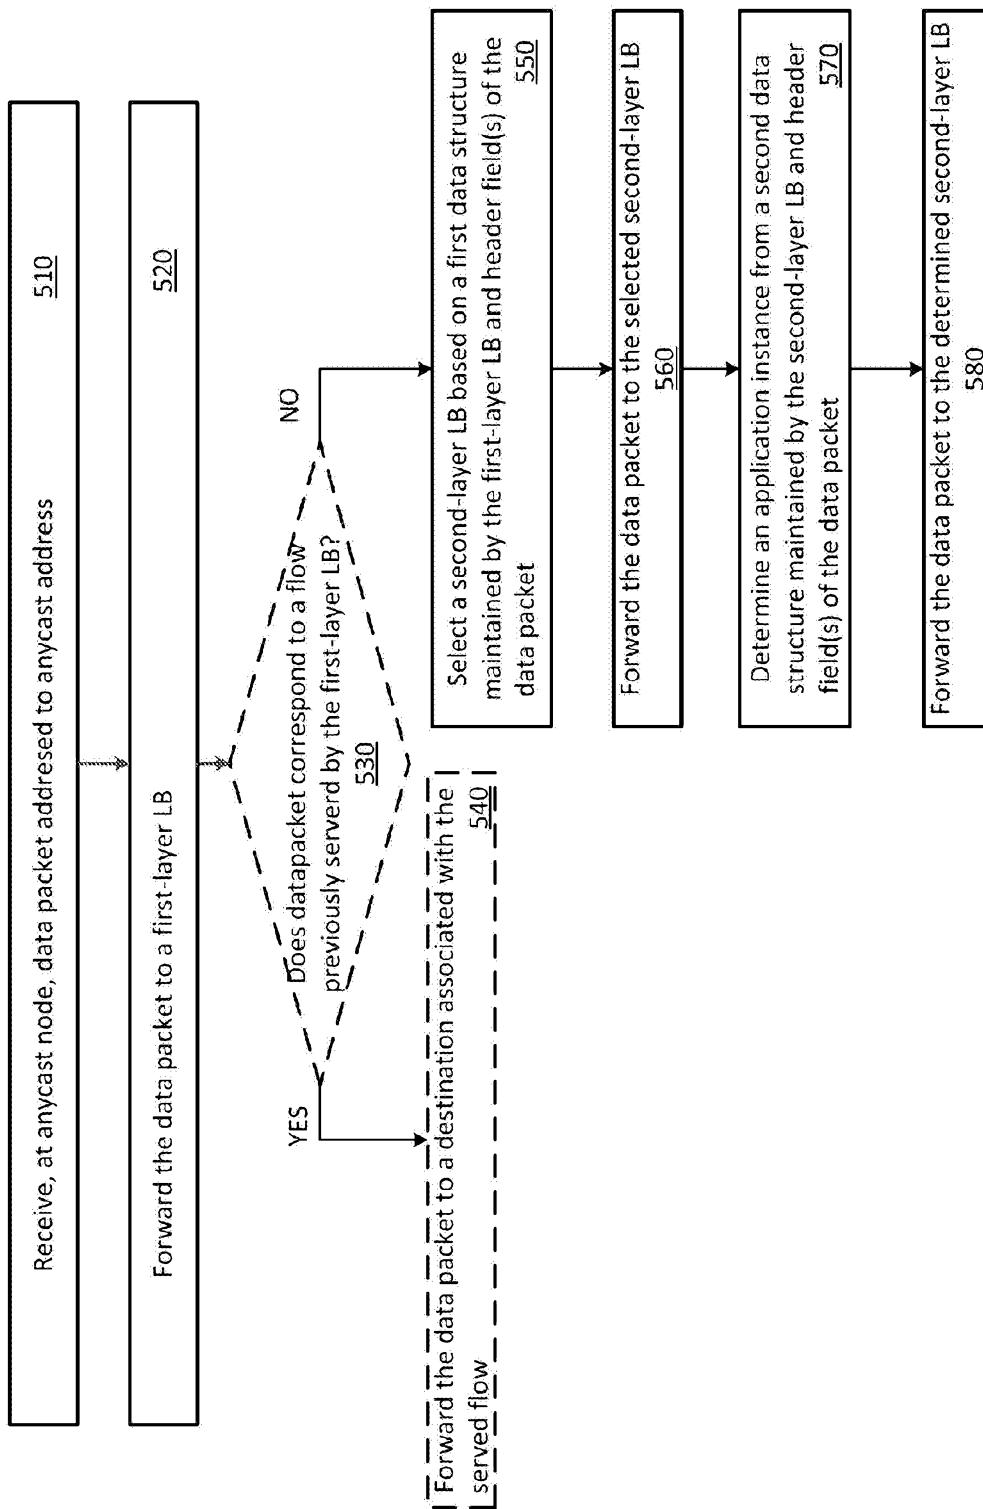


Figure 5

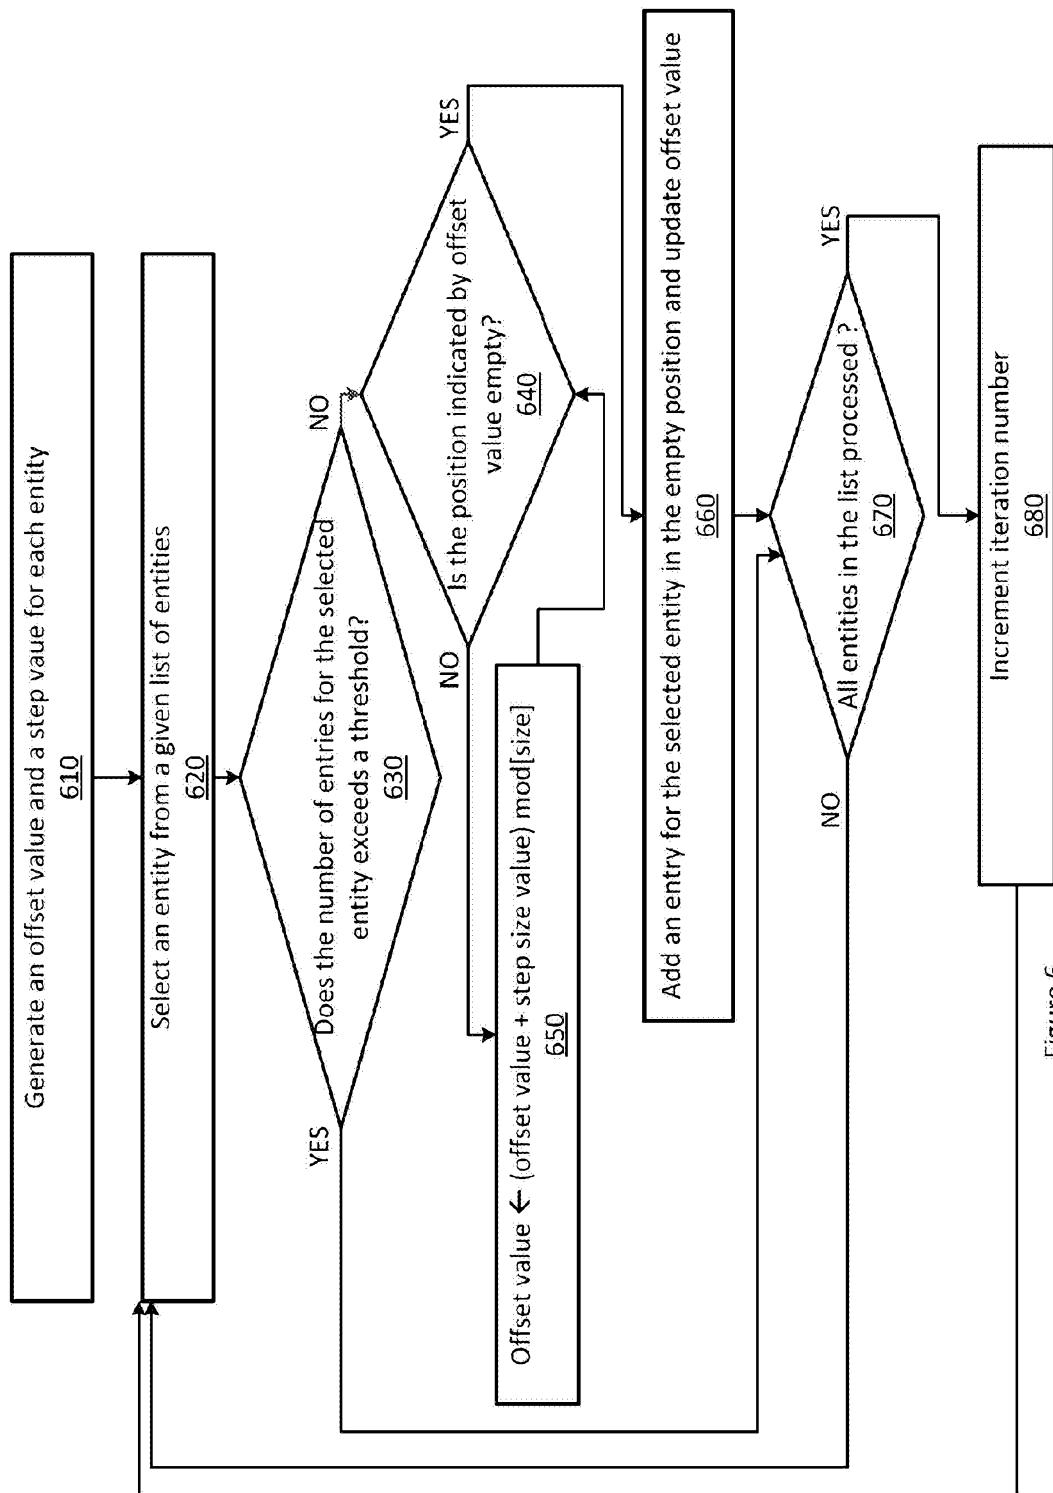


Figure 6

	Weights	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1.	A:1.0; B:0.0	A_2	A_3	A_1	A_1	A_3	A_3	A_1	A_2	A_1	A_2	A_3	A_2	A_1	A_3	A_2	A_2	A_2
2.	A:0.9; B:0.1	A_2	A_3	A_1	A_1	A_3	B_2	A_1	B_3	A_1	A_2	B_1	A_2	A_3	A_2	A_1	A_3	A_2
3.	A:0.8; B:0.2	A_2	A_3	A_1	A_1	A_3	B_2	A_1	B_3	A_1	A_2	B_1	A_2	A_3	A_2	A_1	A_2	A_2
4.	A:0.7; B:0.3	A_2	A_3	B_1	A_1	A_1	A_3	B_2	A_1	B_3	A_1	A_2	B_1	A_2	A_3	A_2	B_2	B_3
5.	A:0.6; B:0.4	A_2	A_3	B_3	A_1	A_1	A_3	B_2	A_1	B_3	B_2	A_2	B_1	A_2	B_1	A_3	A_1	A_2
6.	A:0.5; B:0.5	A_2	A_3	B_3	A_1	B_1	A_3	B_2	A_1	B_3	B_2	A_2	B_1	A_2	B_1	A_3	A_1	B_2
7.	A:0.4; B:0.6	A_2	A_3	B_3	A_1	A_1	A_3	B_2	A_1	B_3	B_2	A_2	B_1	A_2	B_1	B_1	B_2	B_3
8.	A:0.3; B:0.7	B_1	A_3	B_3	A_1	B_2	A_3	B_2	A_1	B_3	B_2	A_2	B_1	A_2	B_1	B_1	B_2	B_3
9.	A:0.2; B:0.8	B_1	B_3	A_1	B_2	A_3	B_2	B_1	B_3	B_2	A_2	B_1	A_2	B_3	A_2	B_2	B_3	
10.	A:0.1; B:0.9	B_1	B_1	B_3	A_1	B_2	A_3	B_2	B_1	B_3	B_2	B_2	B_1	A_2	B_3	B_1	B_2	B_3
11.	A:0.0; B:1.0	B_1	B_1	B_3	B_1	B_2	B_1	B_2	B_3	B_3	B_2	B_2	B_1	B_2	B_3	B_1	B_2	B_3

Figure 7