



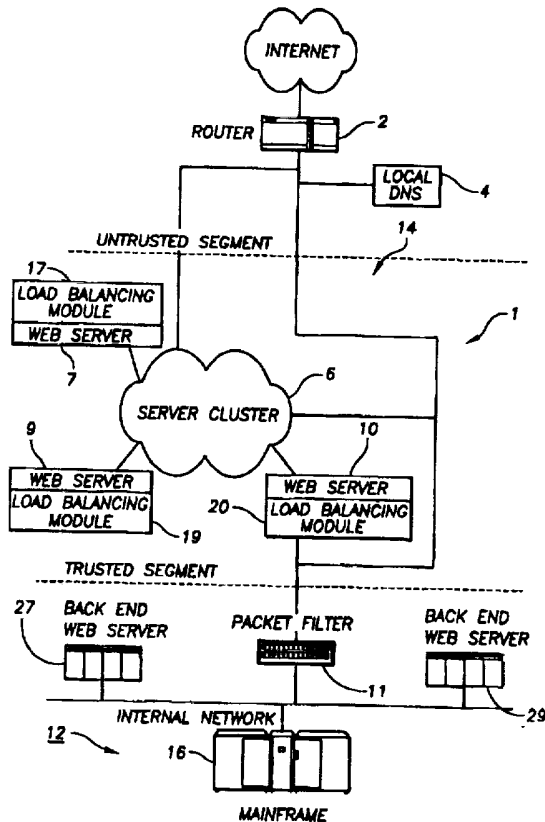
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification<sup>6</sup> : <b>G06K</b></p>	<p><b>A2</b></p>	<p>(11) International Publication Number: <b>WO 99/18534</b> (43) International Publication Date: 15 April 1999 (15.04.99)</p>
<p>(21) International Application Number: PCT/US98/21296 (22) International Filing Date: 6 October 1998 (06.10.98) (30) Priority Data: 60/061,170 6 October 1997 (06.10.97) US 60/071,039 13 January 1998 (13.01.98) US 09/164,499 1 October 1998 (01.10.98) US (71) Applicant: WEB BALANCE, INC. [US/US]; P.O. Box 226, Foxboro, MA 02035-0226 (US). (72) Inventors: O'NEIL, Kevin, M.; 92 West Pond Road, Plymouth, MA 02360 (US). NERZ, Robert, F.; 25 Maryann Way, W. Attleboro, MA 02760 (US). AUBIN, Robert, R.; Apartment 4, 16 Fuller Road, Foxboro, MA 02035 (US). (74) Agent: PYSHER, Paul, A.; Choate, Hall &amp; Stewart, Exchange Place, 53 State Street, Boston, MA 02109 (US).</p>		<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p><b>Published</b> <i>Without international search report and to be republished upon receipt of that report.</i></p>

(54) Title: SYSTEM FOR BALANCING LOADS AMONG NETWORK SERVERS

(57) Abstract

A system which distributes requests among a plurality of network servers receives a request from a remote source at a first one of the network servers, and determines whether to process the request in the first network server. The request is processed in the first network server in a case that it is determined that the request should be processed in the first network server. On the other hand, the request is routed to another network server in a case that it is determined that the request should not be processed in the first network server.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

**SYSTEM FOR BALANCING LOADS AMONG NETWORK SERVERS****Background of the Invention**

5           The present invention is directed to a peer-to-peer load balancing system which is implemented in plural network servers. In particular, the invention is directed to a computer-executable module for use in network servers which enables each server to distribute loads among its peers based on a load currently being processed in each server and/or contents of the network requests. The invention  
10 has particular utility in connection with World Wide Web servers, but can be used with other servers as well, such as CORBA servers, ORB servers, FTP servers, SMTP servers, and Java servers.

          Network systems, such as the World Wide Web (hereinafter "WWW"),  
15 utilize servers to process requests for information. Problems arise, however, if one server becomes overloaded with requests. For example, if a server becomes overloaded, it may be unable to receive new requests, may be slow to process the requests that it has already received, and may yield server errors.

20           Load balancing was developed to address the foregoing problems in the art. Briefly, load balancing involves distributing requests among plural servers (e.g., different servers on a Web site) in order to ensure that any one server does not become unduly burdened.

25           One conventional load balancing technique involves the use of a domain name server (hereinafter "DNS"), in particular a "round-robin" DNS. This device, which typically operates on the network, is responsible for resolving uniform resource locators or "URLs" (e.g., "www.foo.com") to specific IP addresses (e.g., 111.222.111.222). In this regard, a Web site having several servers may operate  
30 under a single URL, although each server is assigned a different IP address. A round-robin DNS performs load balancing by routing requests to these servers in sequential rotation based on their IP addresses.

While round-robin DNSs can coarsely distribute loads among several servers, they have several drawbacks. For example, not all requests for connection to a Web site are necessarily received by a round-robin DNS. Rather, many requests will have been previously "resolved" by a DNS local to the requestor and remote from the Web site (i.e., a "a remote DNS") or by the requestor (i.e., the computer that issued the request on the WWW). In these cases, resolution is based on an address which has been cached in the remote DNS or the requestor, rather than by sequential rotation provided by the Web site's round-robin DNS. Due to this caching, load balancing may not be achieved to a satisfactory degree.

DNS-based load balancing techniques have another significant drawback. In the event that a Web server fails (i.e., the Web server goes off-line), the Web site has no real-time mechanism by which to reroute requests directed to that server (e.g., by a remote DNS). Thus, a remote DNS with caching capabilities could continue to route requests to a failed server for hours, or even days, after the failure has occurred. As a result, a user's connection would be denied with no meaningful error message or recovery mechanism. This situation is unacceptable, particularly for commercial Web sites.

As an alternative to the DNS-based load balancing techniques described above, some vendors have introduced dedicated load balancing hardware devices into their systems. One such system includes a device, called a proxy gateway, which receives all network requests and routes those requests to appropriate Web servers. In particular, the proxy gateway queries the servers to determine their respective loads and distributes network requests accordingly. Responses from the servers are routed back to the network through the proxy gateway. Unlike the DNS-based schemes, all requests resolve to the IP address of the proxy server, thereby avoiding the risk that remote DNS caching or failed servers will inadvertently thwart access to the site.

While proxy gateways address some of the fundamental problems of load balancing described above, they also have several drawbacks. For example, proxy

gateways add latency in both the "request" direction and the "response" direction. Moreover, since the proxy gateway is, for all intents and purposes, the only way into or out of a Web site, it can become a bottleneck that limits the capacity of that site to the capacity of the proxy gateway. Moreover, the proxy gateway is also a  
5 single point of failure, since its failure alone will prevent access to the Web site.

An IP redirector is a device similar to a proxy gateway which also performs load balancing. Like a proxy gateway, an IP redirector serves as a hub that receives and routes requests to appropriate servers based on the servers' loads. IP  
10 redirectors are different from proxy gateways in that IP redirectors do not handle responses to requests, but rather let those responses pass directly from the assigned Web servers to the requesters. However, IP redirectors suffer from many of the same drawbacks of the proxy gateways described above, particularly insofar as limiting the capacity of the Web site and preventing access to it as a result of  
15 failure of the IP redirector.

Dedicated load balancers, such as proxy gateways and IP redirectors, also have drawbacks related to sensing loads in different Web servers. Using current technologies, a server can become busy in a matter of milliseconds. A load  
20 balancer, however, can only query various servers so often without creating undesirable overhead on the network and in the servers themselves. As a result, such load balancers often must rely on "old" information" to make load balancing decisions. Load balancing techniques which use this "old" information are often ineffective, particularly in cases where such information has changed significantly.  
25

Dedicated load balancers, such as proxy gateways and IP redirectors, also have problems when it comes to electronic commerce transactions. In this regard, electronic commerce transactions are characterized by multiple sequential requests from a single client, where each subsequent request may need to refer to state  
30 information provided in an earlier request. Examples of this state information include passwords, credit card numbers, and purchase selections.

Problems with electronic commerce arise because an entire transaction must be serviced by one of plural network servers, since only that one server has the original state information. A load balancer therefore must identify the first request of a stateful transaction and keep routing requests from that requestor to the same server for the duration of the transaction. However, it is impossible for the load balancer to know exactly where a transaction begins or ends, since the information in the request providing such indications may be encrypted (e.g., scrambled) when it passes through the dedicated load balancer. In order to maintain an association between one requestor and one server, dedicated load balancers therefore use a mechanism referred to as a "sticky timer". More specifically, the load balancer infers which request may be the start of a stateful transaction and then sets a "sticky timer" of arbitrary duration (e.g., 20 minutes) which routes all subsequent requests from the same requestor to the same Web server, and which renews the "sticky timer" with each subsequent request. This method is easily bypassed and may unnecessarily defeat the load balancing feature.

Thus, there exists a need for a load balancing technique which is able to provide more accurate load balancing than the techniques described above, which is able to perform accurate load balancing despite cached server addresses or "maintained" Web browser addresses, which is not a significant bottleneck or source of single point failure, and which is able to maintain the association between a client and a server in order to preserve state information required to complete an electronic commerce transaction.

## 25 **Summary of the Invention**

The present invention addresses the foregoing needs by providing, in one aspect, a plurality of network servers which directly handle load balancing on a peer-to-peer basis. Thus, when any of the servers receives a request, the server either processes the request or routes the request to one of its peers -- depending on their respective loads and/or on the contents of the request. By implementing load balancing directly on the servers, the need for dedicated load balancing hardware is

reduced, as are the disadvantages resulting from such hardware. Thus, for example, because each server has the capability to perform load balancing, access to a Web site managed by the server is not subject to a single point of failure. Moreover, requests tagged with IP addresses cached by remote DNSs or the requestor itself  
5 are handled in the same way as other requests, i.e., by being routed among the load balancing-enabled servers.

A network server according to a related aspect of the invention exchanges information with its peers regarding their respective loads. This exchange may be  
10 implemented based on either a query/response or unsolicited multicasts among the server's peers, and may be encrypted or may occur over a private communication channel. The exchange may be implemented to occur periodically or may be triggered by a network event such as an incoming request. In a preferred embodiment of the invention, each server multicasts its load information to its  
15 peers at a regular period (e.g., 500 ms). This period may be set in advance and subsequently re-set by a user. In the preferred embodiment, the multicast message serves the dual purposes of exchanging load information and of confirming that a transmitting server is still on-line.

By virtue of the foregoing, and by virtue of the server having nearly  
20 instantaneous information regarding its own workload, the server is able to make routing determinations based on substantially up-to-date information. The most critical decision, i.e., whether to consider rerouting, is preferably made based on the most current information available (i.e., based on a local server load provided  
25 nearly instantaneously from within the server and without any network transmission latency).

In further aspects of the invention, a server processes a received request directly when its load is below a first predetermined level, or if its load is above  
30 the first predetermined level yet those of the server's peers are above a second predetermined level. Otherwise, the server routes the request to one of its peers. By equipping a site with multiples servers of this type, it is possible to reduce the

chances that one server will become overwhelmed with requests while another server of similar or identical capabilities remains relatively idle.

5 In other aspects of the invention, the receiving server determines whether to process a request based on its content, e.g., its uniform resource indicator ("URI"). By virtue of this feature of the invention, it is possible to limit the server to processing certain types of network requests, while routing others. Alternatively, it is possible to direct particular requests to particular servers, which then may either process or re-route those requests based on loads currently being handled by the  
10 servers.

In other aspects of the invention, the receiving server determines which, if any, of its peer servers are off-line. The server then routes requests to its on-line peers and does not route requests to its off-line peers. A server may also assume  
15 the network identity (i.e., the IP address and/or URL) of an off-line peer to insure that requests are serviced properly even if directed to an off-line peer by virtue of caching in a remote DNS. The server would continue to service both its own identity and its assumed identity until the off-line peer returns to on-line service. As a result, it is possible to reduce response errors resulting from requests being  
20 inadvertently directed to off-line servers.

In other aspects of the invention, servers may be configured to recognize specific URIs which designate entry points for stateful transactions. A server so configured will not re-route requests away from itself if they are related to a  
25 stateful transaction conforming to the URI of the server. Even URIs that arrive in encrypted requests will be decrypted by the server and, therefore, will be subject to intelligent interpretation in accordance with configuration rules. As a result, an electronic commerce transaction comprised of multiple requests may be processed entirely on one of plural servers. Once the transaction is complete, as confirmed  
30 by comparing URI information to configuration rules, subsequent requests will again be subject to re-routing for the purpose of load balancing.



This brief summary has been provided so that the nature of the invention may be understood quickly. A more complete understanding of the invention can be obtained by reference to the following detailed description of the preferred embodiments thereof in connection with the attached drawings.

5

### **Brief Description of the Drawings**

A more complete understanding of the invention may be attained by reference to the drawings, in which:

10

Figure 1 is a diagram showing the topology of a Web site including the present invention;

Figure 2, comprised of Figures 2A and 2B, is a flow diagram showing process steps for distributing requests among various servers based on the loads being handled by them;

15

Figure 3 is a more detailed view of a portion of the topology shown in Figure 1 relating to load balancing;

20

Figure 4 is a flow diagram showing process steps for distributing requests among various servers based on the content of the requests; and

Figure 5 is a diagram showing the topology of a Web site including the present invention and a proxy.

25

### **Detailed Description of the Illustrated Embodiments**

The present invention is directed to a system for implementing peer-to-peer load balancing among plural network servers. Although the invention will be described in the context of the World Wide Web ("WWW"), and more specifically in the context of WWW servers, it is not limited to use in this context. Rather, the

30

invention can be used in a variety of different types of network systems with a variety of different types of servers. For example, the invention can be used in intranets and local area networks, and with CORBA servers, ORB servers, FTP servers, SMTP servers, and Java servers, to name a few.

5

Figure 1 depicts the topology of a Web site 1 which includes the present invention, together with hardware for accessing that Web site from a remote location on the Internet. More specifically, Figure 1 shows router 2, local DNS 4, server cluster 6 comprised of Web servers 7, 9 and 10, packet filter 11, and internal network 12. A brief description of this hardware is provided below.

Router 2 receives requests for information stored on Web site 1 from a remote location (not shown) on the Internet. Router 2 routes these requests, which typically comprise URLs, to local DNS 4. Local DNS 4 receives a URL from router 2 and resolves the domain name in the URL to a specific IP address in server cluster 6.

Server cluster 6 is part of the untrusted segment 14 of Web site 1, to which access is relatively unrestricted. Server cluster 6 is comprised of a plurality of servers, including servers 7, 9 and 10. Each of these servers is capable of retrieving information from internal network 12 in response to requests resolved by a remote DNS on the Internet or by local DNS 4. Included on each of servers 7, 9 and 10 is a microprocessor (not shown) and a memory (not shown) which stores process steps to effect information retrieval. In preferred embodiments of the invention, each memory is capable of storing and maintaining programs and other data between power cycles, and is capable of being reprogrammed periodically. An example of such a memory is a rotating hard disk.

The memory on each server also stores a computer-executable module (i.e., a heuristic) comprised of process steps for performing the peer-to-peer load balancing technique of present invention. More specifically, server 7 includes load balancing module 17, server 9 includes load balancing module 19, and server 10

includes load balancing module 20. The process steps in these modules are executable by the microprocessor on each server so as to distribute requests among the Web servers. In more detail, the process steps include, among other things, code to receive a request from a remote source at a first one of the Web servers  
5 (e.g., server 7), code to determine whether to process the request in the first server, code to process the request in the first server in a case that the determining code determines that the request should be processed in the first server, and code to route the request to another server (e.g., server 9) in a case that the determining code determines that the request should not be processed in the first server. A  
10 more detailed description of the load balancing technique implemented by these process steps is provided below.

Packet filter 11 comprises a firewall for internal network 12 (i.e., the trusted segment) of Web site 1. All transactions into or out of internal network 12 are  
15 conducted through packet filter 11. In this regard, packet filter 11 "knows" which inside services of internal network 12 may be accessed from the Internet, which clients are permitted access to those inside services, and which outside services may be accessed by anyone on internal network 12. Using this information, packet filter 11 analyzes data packets passing therethrough and filters these packets accordingly,  
20 restricting access where necessary and allowing access where appropriate.

Internal network 12 includes mainframe 16 and back-end Web servers 27 and 29. Back-end Web servers 27 and 29 comprise file servers which store a database for Web site 1. Back-end Web servers 27 and 29 may be used to access  
25 data files on mainframe 16 (or other similar computer) in response to requests from server cluster 6. Once such data files have been accessed, mainframe 16 may then transmit these files back to server cluster 6. Alternatively, data on back-end Web servers 27 and 29 may be accessed directly from server cluster 6 without the aid of mainframe 16.

30

#### First Embodiment

Figure 2 illustrates process steps of the present invention for load balancing`

received network requests. To begin, in step S201 a network request is received at a server, such as server 7 shown in Figure 3. This request may be resolved by a remote DNS on the Internet based on a cached IP address (e.g., requests 1, 2, 3 and 4) or, alternatively, the request may be resolved by a local round-robin DNS 4 (e.g., request 5). Then, in step S202, server 7 determines a load (e.g., the number and/or complexity of network requests) that it is currently processing, and the capacity remaining therein.

Step S203 decides if the load currently being processed in server 7 exceeds a first predetermined level. In preferred embodiments of the invention, this predetermined level is 50%, meaning that server 7 is operating at 50% capacity. Of course, the invention is not limited to using 50% as the first predetermined level. In this regard, a value for the first predetermined level may be stored in a memory on server 7, and may be reprogrammed periodically.

15

If step S203 decides that server 7 is not processing a load that exceeds the first predetermined level, flow proceeds to step S204. In step S204, the network request is processed in server 7, and a response thereto is output via the appropriate channels. On the other hand, in a case that step S203 determines that server 7 is processing a load that exceeds the first predetermined level, flow proceeds to step S205.

20

Step S205 determines loads currently being processed by server 7's peers (e.g., servers 9 and 10 shown in Figure 3). In more detail, in step S205, load balancing module 17 compares its current load information with the most recent load information provided by load balancing modules 19 and 20. These load balancing modules continuously exchange information regarding their respective loads, so that this information is instantly available for comparison. In the example shown in Figure 3, load balancing module 19 provides information concerning the load currently being processed by server 9, and load balancing module 20 provides information concerning the load currently being processed by server 10.

30

In step S206, load balancing module 17 determines whether the loads currently being processed by server 7's peers are less than the load on server 7 by a differential exceeding a second predetermined level. In preferred embodiments of the invention, this second predetermined level is 20%, which provides a means of assessing whether servers 9 or 10 have at least 20% more of their capacities available than server 7. Of course, the invention is not limited to using 20% as the second predetermined level. In this regard, as above, a value for the second predetermined level may be stored in a memory on server 7, and may be reprogrammed periodically.

In a case that step S206 decides that server 7's peers (i.e., servers 9 and 10) do not have 20% more of their capacity available, flow proceeds to step S204. In step S204, the network request is processed in server 7, and a response thereto is output via the appropriate channels. On the other hand, in a case that step S206 decides that at least one of server 7's peers is processing a load that is less than the percent load on server 7 by the second predetermined level, flow proceeds to step S207.

Step S207 determines which, if any, of the servers at Web site 1 are off-line based, e.g., on the load information exchange (or lack thereof) in step S205. A server may be off-line for a number of reasons. For example, the server may be powered-down, malfunctioning, etc. In such cases, the servers' load balancing modules may be unable to respond to a request from load balancing module 17 or otherwise be unable to participate in an exchange of information, thereby indicating that those servers are off-line. In addition, in preferred embodiments of the invention, the load balancing modules are able to perform diagnostics on their respective servers. Such diagnostics test operation of the servers. In a case that a server is not operating properly, the server's load balancing module may provide an indication to load balancing module 17 that network requests should not be routed to that server.

Next, step S208 analyzes load information from on-line servers in order to

determine which of the on-line servers is processing the smallest load. Step S208 does this by comparing the various loads being processed by other servers 9 and 10 (assuming that both are on-line). Step S209 then routes the network request to the server which is currently processing the smallest load. In the invention, routing is performed by sending a command from load balancing module 17 to a requestor instructing the requestor to send the request to a designated server. Thus, re-routing is processed automatically by the requestor software and is virtually invisible to the actual Internet user.

10           Thereafter, that server processes the request in step S210. At this point, it is noted, however, that the invention is not limited to routing the request to a server that is processing the smallest load. Rather, the invention can be configured to route the request to any server that is operating at or below or predetermined capacity, or something similar such as, but not limited to, a round-robin hand-off rotation.

15           Figure 3 illustrates load distribution according to the present invention. More specifically, as noted above, server 7 (more specifically, load balancing module 17) receives requests 1, 2, 3 and 4 resolved by network DNS 21 and request 5 via local DNS 4. Similarly, server 10 receives request 6 (i.e., a cached request) via local DNS 4. Any of these requests may be "bookmarked" requests, meaning that they are specifically addressed to one server. Once each load balancing module receives a request, it determines whether to process that request in its associated server or to route that request to another server. This is done in the manner shown in Figure 2. By virtue of the processing shown in Figure 2, load balancing modules 17, 19 and 20 distribute requests so that server 7 processes requests 1 and 2, server 9 processes requests 3 and 5, and server 10 processes requests 4 and 6.

### 30   Second Embodiment

          In the second embodiment of the invention, load balancing is performed based on a content of a network request, in this case a URL/URI. As noted above,

a URL addresses a particular Web site and takes the form of "www.foo.com". A URI, on the other hand, specifies information of interest at the Web site addressed by the URL. For example, in a request such as "www.foo.com/banking", "/banking" is the URI and indicates that the request is directed to information at the "foo" Web site that relates to "banking". In this embodiment of the invention, URIs in network requests are used to distribute requests among servers.

Figure 4 is a flow diagram illustrating process steps comprising this embodiment of the invention. To begin, in step S401, load balancing module 17 receives a request from either network DNS 21 or from local DNS 4 (see Figure 3). In step S402, the load balancing module then analyzes the request to determine its content. In particular, load balancing module 17 analyzes the request to identify URIs (or lack thereof) in the request.

Step S402 determines which server(s) are dedicated to processing which URIs, and which server(s) are dedicated to processing requests having no URI. That is, in the invention, the load processing module of each server is configured to accept requests for one or more URIs, thus limiting the server to processing requests for those URIs. For example, load balancing module 17 may be configured to accept requests with a URI of "/banking", whereas load balancing module 19 may be configured to accept requests with a URI of "/securities". Which server processes which URI may be "hard-coded" within the server's loading balancing module, stored within the memory of each server, or obtained and updated via a dynamic protocol.

In any event, in a case that step S403 decides that server 17 is dedicated to processing URIs of the type contained in the request (or no URI, whichever the case may be), flow proceeds to step S404. In step S404, the request is accepted by load balancing module 17 and processed in server 7, whereafter processing ends. On the other hand, in a case that step S403 decides that server 7 does not process URIs of the type contained in the request, flow proceeds to step S405. This step routes the request to one of server 7's peers that is dedicated to processing requests

containing such URIs. Routing is performed in the same manner as in step S209 of Figure 2. Once the request is received at the appropriate server, the load balancing module associated therewith accepts the request for processing by the server in step S406, whereafter processing ends.

5

### Third Embodiment

The first and second embodiments of the invention described above can be combined into a single embodiment which routes network requests based on both a content of the request and loads being handled by the various servers. More specifically, in this embodiment of the invention, each load balancing module is configured to route a request to a server dedicated to a particular URI in a case that the server is operating at less than a predetermined capacity. In a case that the server is operating at above the predetermined capacity, the invention routes the requests to another server which can handle requests for the URI, but which is operating at below the predetermined capacity. The methods for performing such routing are described above with respect to the first and second embodiments of the invention.

### Fourth Embodiment

As noted above, the present invention reduces the need for a proxy gateway or similar hardware for distributing loads among various Web servers. It is noted, however, that the invention can be used with such hardware. Figure 5 shows the topology of a Web site on which the present invention is implemented, which also includes proxy 26.

25

In this regard, except for proxy 26, the features show in Figure 5 are identical in both structure and function to those shown in Figure 1. With respect to proxy 26, proxy 26 is used to receive network requests and to route those requests to appropriate servers. A load balancing module on each server then determines whether the server can process requests routed by proxy 26 or whether such requests should be routed to one of its peers. The process for doing this is set forth in the first, second and third embodiments described above.

30



Fifth Embodiment

This embodiment of the invention is directed to a system for maintaining an association between a requestor and one of plural servers at a Web site when state information is used during an electronic transaction.

5

More specifically, in accordance with this embodiment of the invention, a server at a Web site, such as server 7 shown in Figure 1, is configured to recognize specific URIs (e.g., URIs that designate entry points for a stateful transaction relating to electronic commerce). In the case that one of these URIs is recognized, the server will not route subsequent transactions away from that server, thereby ensuring that all such requests are processed by that server. Requests may again be re-routed from the server once a URI which matches a predetermined "configuration rule" is detected (e.g., when a transaction is complete).

15

In preferred embodiments of the invention, wild card URI information may be used to designate a stateful path. For example, the hyperlink "http://www.foo.com/banking/\*" would mean that "http://www.foo.com/banking/" constitutes the entry point into a stateful transaction. Any request up to and including this point would be subject to potential re-routing. Any request further down this path would indicate that the requestor and the server are engaged in a stateful transaction and not subject to potential re-routing.

20

The present invention has been described with respect to particular illustrative embodiments. It is to be understood that the invention is not limited to the above-described embodiments and modifications thereto, and that various changes and modifications may be made by those of ordinary skill in the art without departing from the spirit and scope of the appended claims.

25

In view of the foregoing, what we claim is:

30

1. A method of distributing requests among a plurality of network servers, the method comprising the steps of:

5 receiving a request from a remote source at a first one of the network servers;

executing a determining step in the first server, the determining step for determining whether to process the request in the first network server;

10 processing the request in the first network server in a case that the determining step determines that the request should be processed in the first network server; and

15 routing the request to another network server in a case that the determining step determines that the request should not be processed in the first network server.

2. A method according to claim 1, wherein the determining step makes a determination as to whether the request should be processed in the first network server based on a load currently being processed in the first network server.

20

3. A method according to claim 2, wherein the determining step makes the determination based, in addition, on a load currently being processed in one or more of the other network servers.

25 4. A method according to claim 1, wherein the determining step comprises the steps of:

determining a load currently being processed by the first network server;  
and

30

receiving information in the first network server from each of the other network servers, the information from each of the other network servers comprising

information concerning a load currently being processed in each network server;

5 wherein the determining step determines that the first network server should process the request in a case that (i) the load currently being processed in the first network server is below a first predetermined level, or (ii) the load currently being processed in the first network server is above the first predetermined level and is above loads currently being processed by either of the other network servers by less than a second predetermined level; and

10 wherein the determining step determines that the first network server should not process the request in a case that the load currently being processed in the first network server is above the first predetermined level and a load currently being processed in at least one of the other network servers is below the level of the first network server by at least the second predetermined level.

15

5. A method according to claim 1, wherein, in a case that the determining step determines that the request should not be processed in the first network server and the plurality of network servers includes at least two other network servers, the method further comprises a second determining step for determining which of the at least two other network servers that the request should be routed to in the routing step.

20

6. A method according to claim 5, wherein the second determining step determines which of the at least two other network servers that the request should be routed to based on loads currently being processed in the at least two other network servers.

25

7. A method according to claim 6, wherein the second determining step determines that the request should be routed to a network server which is currently processing a smallest load.

30

8. A method according to claim 1, wherein the plurality of network servers

comprises one or more of the following types of servers: World Wide Web servers, CORBA servers, ORB servers, FTP servers, and SMTP servers.

9. A method according to claim 1, wherein the routing step comprises  
5 sending a command to the remote source which instructs the remote source to send the request to the other one of the network servers.

10. A method according to claim 1, wherein the determining step  
10 determines whether to process the request in the first network server based on a content of the request.

11. A method according to claim 10, wherein the request comprises a  
uniform resource locator ("URL") and a uniform resource indicator ("URI"); and  
15 wherein the determining step determines whether to process the request in the first network server based on the URI in the request.

12. A method according to claim 11, wherein the determining step  
determines whether to process the request in the first network server based, in  
20 addition, on a load currently being processed in the first network server and loads currently being processed in one or more of the other network servers.

13. A method according to claim 1, further comprising, before the routing  
step, the step of determining which, if any, of the plurality of network servers are  
25 off-line;

wherein the routing step routes the request to a network server which is on-line and does not route the request to a network server which is off-line.

30 14. Computer-executable process steps stored on a computer-readable medium, the computer executable process steps comprising a server module which is installable in a plurality of network servers to distribute requests among the

plurality of network servers, the computer-executable process steps comprising:

code to receive a request from a remote source at a first one of the network servers;

5

code, executable by the first server, to determine whether to process the request in that server;

code to process the request in the first network server in a case that the determining code determines that the request should be processed in the first network server; and

10

code to route the request to another network server in a case that the determining code determines that the request should not be processed in the first network server.

15

15. Computer-executable process steps according to claim 14, wherein the determining code comprises code to make a determination as to whether the request should be processed in the first network server based on a load currently being processed in the first network server.

20

16. Computer-executable process steps according to claim 15, wherein the determining code comprises code to make the determination based, in addition, on a load currently being processed in one or more of the other network servers.

25

17. Computer-executable process steps according to claim 14, wherein the determining code comprises:

code to determine a load currently being processed by the first network server; and

30

code to receive information in the first network server from each of the

other network servers, the information from each of the other network servers comprising information concerning a load currently being processed in each network server;

5            wherein the determining code determines that the first network server should process the request in a case that (i) the load currently being processed in the first network server is below a first predetermined level, or (ii) the load currently being processed in the first network server is above the first predetermined level and is above loads currently being processed by either of the other network servers by less  
10 than a second predetermined level; and

             wherein the determining code determines that the first network server should not process the request in a case that the load currently being processed in the first network server is above the first predetermined level and a load currently being  
15 processed in at least one of the other network servers is below the level of the first network server by at least the second predetermined level.

             18. Computer-executable process steps according to claim 14, wherein the computer-executable process steps further comprise code to determine which of the  
20 other network servers that the request should be routed to by the routing code in a case that the determining code determines that the request should not be processed in the first network server.

             19. Computer-executable process steps according to claim 18, wherein the  
25 code to determine which of the at least two other network servers that the request should be routed to makes a determination based on loads currently being processed in the at least two other network servers.

             20. Computer-executable process steps according to claim 19, wherein the  
30 code to determine which of the at least two other network servers that the request should be routed to determines that the request should be routed to a network server which is currently processing a smallest load.

21. Computer-executable process steps according to claim 14, wherein the plurality of network servers comprises one or more of the following types of servers: World Wide Web servers, CORBA servers, ORB servers, FTP servers, and SMTP servers.

5

22. Computer-executable process steps according to claim 14, wherein the routing code comprises code to send a command to the remote source which instructs the remote source to send the request to the other one of the network servers.

10

23. Computer-executable process steps according to claim 14, wherein the determining code determines whether to process the request in the first network server based on a content of the request.

15

24. Computer-executable process steps according to claim 23, wherein the request comprises a uniform resource locator ("URL") and a uniform resource indicator ("URI"); and

wherein the determining code determines whether to process the request in the first network server based on the URI in the request.

20

25. Computer-executable process steps according to claim 24, wherein the determining code determines whether to process the request in the first network server based, in addition, on a load currently being processed in the first network server and loads currently being processed in one or more of the other network servers.

25

26. Computer-executable process steps according to claim 14, further comprising code to determine which, if any, of the plurality of network servers are off-line;

30

wherein the routing code routes the request to a network server which is on-

line and does not route the request to a network server which is off-line.

27. A network server which is capable of processing requests and of distributing the requests among a plurality of other network servers, the network  
5 server comprising:

a memory which stores a module comprised of computer-executable process steps; and

10 a processor which executes the process steps stored in the memory so as (i) to receive a request from a remote source at the network server, (ii) to determine whether to process the request in the network server, (iii) to process the request in the network server in a case that the processor determines that the request should be processed in the network server, and (iv) to route the request to another one of  
15 the plurality of network servers in a case that the processor determines that the request should not be processed in the network server.

28. A network server according to claim 27, wherein the processor makes a determination as to whether the request should be processed in the network server  
20 based on a load currently being processed in the network server.

29. A network server according to claim 27, wherein the processor makes the determination based, in addition, on a load currently being processed in one or more of the other network servers.

25 30. A network server according to claim 27, wherein the processor determines whether to process the request in the network server by executing process steps so as (i) to determine a load currently being processed by the first network server, and (ii) to receive information in the first network server from each  
30 of the other network servers, the information from each of the other network servers comprising information concerning a load currently being processed in each network server;



wherein the processor determines that the first network server should process the request in a case that (i) the load currently being processed in the first network server is below a first predetermined level, or (ii) the load currently being processed in the first network server is above the first predetermined level and is  
5 above loads currently being processed by either of the other network servers by less than a second predetermined level; and

wherein the processor determines that the first network server should not process the request in a case that the load currently being processed in the first  
10 network server is above the first predetermined level and a load currently being processed in at least one of the other network servers is below the level of the first network server by at least the second predetermined level.

31. A network server according to claim 27, wherein, in a case that the  
15 processor determines that the request should not be processed in the network server and the plurality of other network servers includes at least two other network servers, the processor executes process steps to determine to which of the at least two other network servers that the request should be routed.

20 32. A network server according to claim 31, wherein the processor determines which of the at least two other network servers that the request should be routed to based on loads currently being processed in the at least two other network servers.

25 33. A network server according to claim 32, wherein the processor determines that the request should be routed to a network server which is currently processing a smallest load.

30 34. A network server according to claim 27, wherein the plurality of other network servers comprises one or more of the following types of servers: World Wide Web servers, CORBA servers, ORB servers, FTP servers, and SMTP servers.

35. A network server according to claim 27, wherein the processor routes the request to another one of the plurality of network servers by executing process steps to send a command to the remote source which instructs the remote source to send the request to the other one of the network servers.

5

36. A network server according to claim 27, wherein the processor determines whether to process the request in the network server based on a content of the request.

10

37. A network server according to claim 36, wherein the request comprises a uniform resource locator ("URL") and a uniform resource indicator ("URI"); and

wherein the processor determines whether to process the request in the network server based on the URI in the request.

15

38. A network server according to claim 37, wherein the processor determines whether to process the request in the network server based, in addition, on a load currently being processed in the network server and a load currently being processed in one or more of the other network servers.

20

39. A network server according to claim 27, wherein the processor executes process steps to determine which, if any, of the plurality of network servers are off-line;

25

wherein the processor routes the request to a network server which is on-line and does not route the request to a network server which is off-line.

40. A method according to Claim 1, wherein the determining step comprises determining whether the request is related to a stateful transaction based on a URI in the request; and

30

wherein (i) in a case that the request is related to a stateful transaction,

determining that the request should be processed in the first network server, and (ii) in a case that the request is not related to a stateful transaction, determining if the request should be processed in the first network server.

5           41. A method according to Claim 40, wherein, in a case that the request is related to a stateful transaction, determining that at least a second request having a URI substantially the same as the URI of the request should be processed in the first network server.

10           42. Computer-executable process steps according to Claim 14, wherein the determining code comprises code to determine whether the request is related to a stateful transaction based on a URI in the request; and

              wherein (i) in a case that the request is related to a stateful transaction, the  
15           determining code determines that the request should be processed in the first network server, and (ii) in a case that the request is not related to a stateful transaction, the determining code determines if the request should be processed in the first network server.

20           43. Computer-executable process steps according to Claim 42, wherein, in a case that the request is related to a stateful transaction, the code to determine determines that at least a second request having a URI substantially the same as the URI of the request should be processed in the first network server.

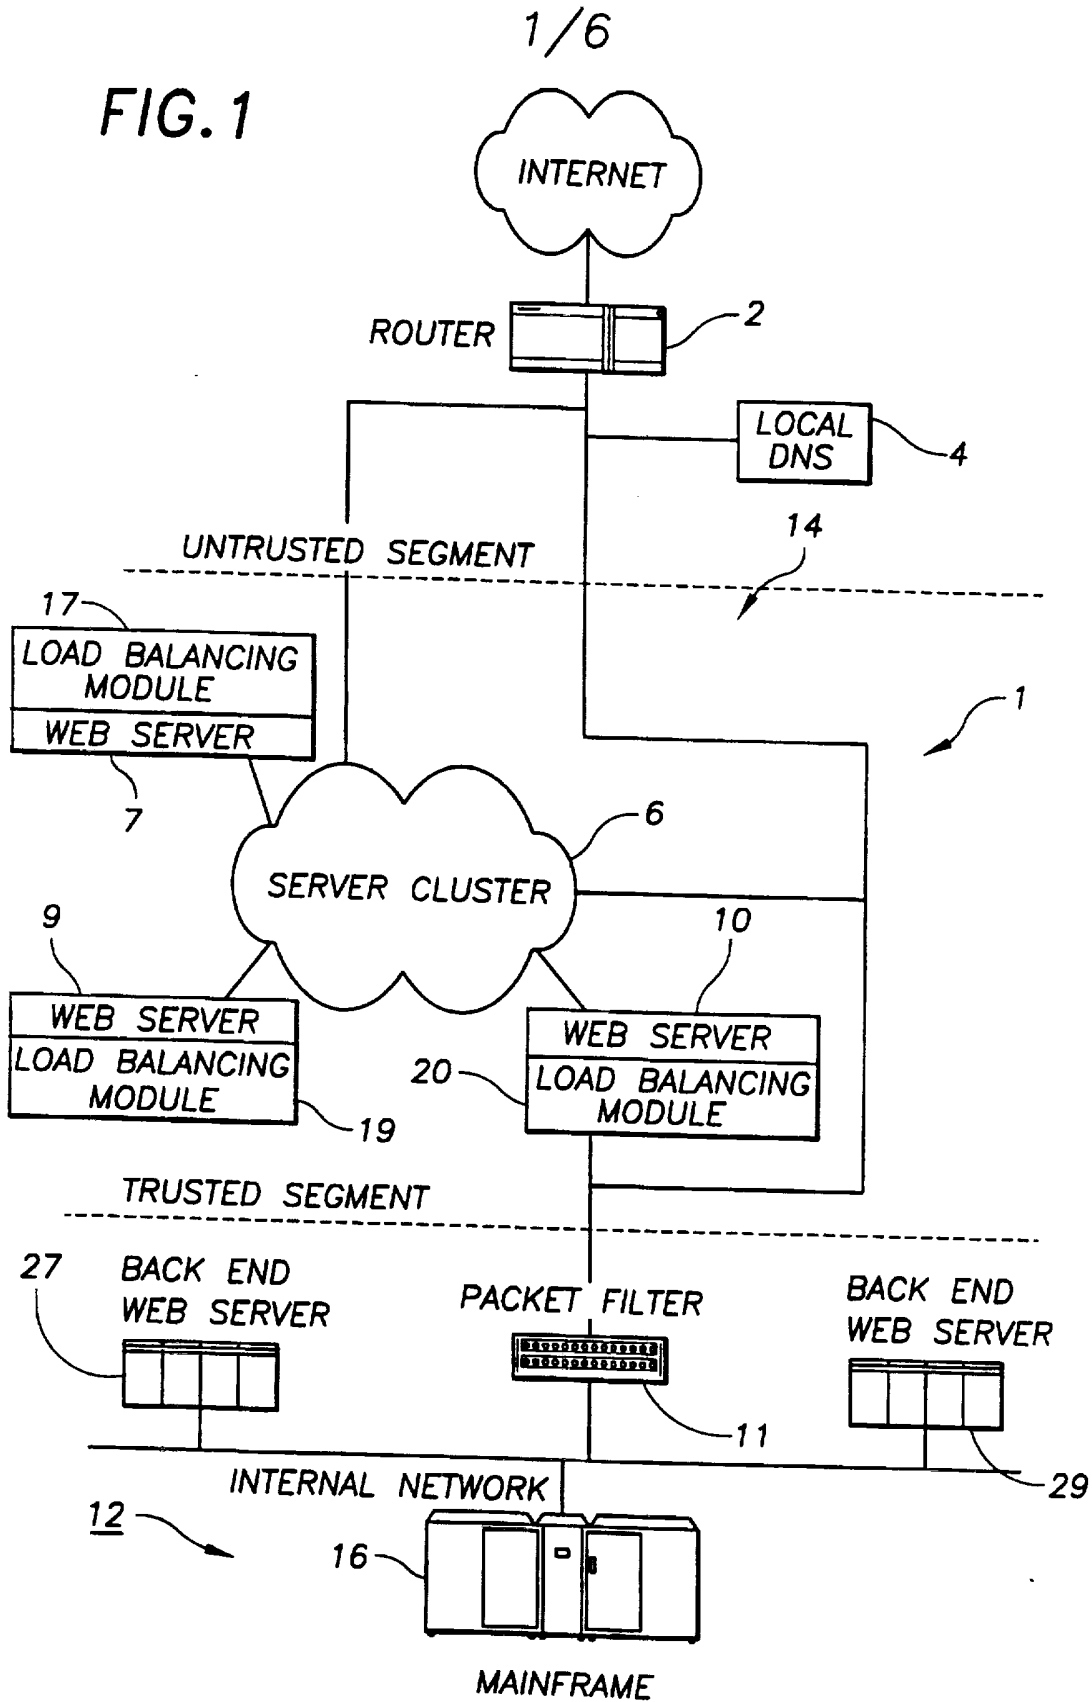
25           44. A network server according to Claim 27, wherein the processor determines whether the request should be processed in the network server by determining whether the request is related to a stateful transaction based on a URI in the request; and

30           wherein (i) in a case that the request is related to a stateful transaction, the processor determines that the request should be processed in the network server, and (ii) in a case that the request is not related to a stateful transaction, the

processor determines if the request should be processed in the network server.

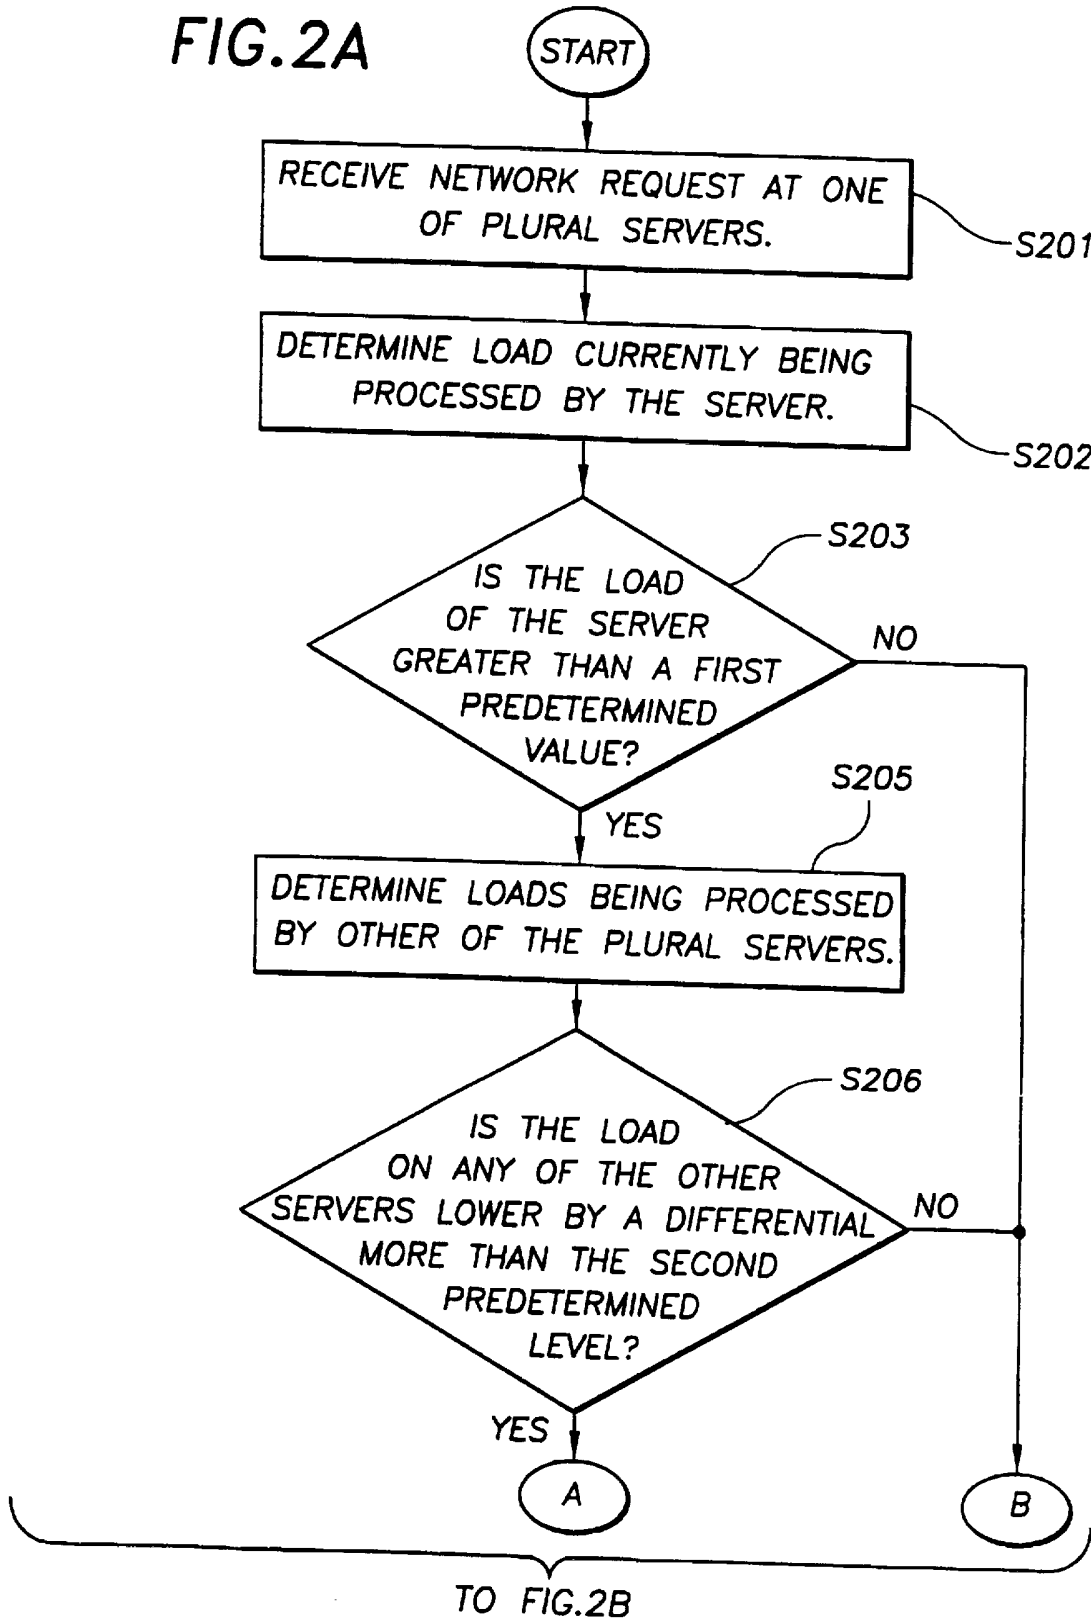
45. A network server according to Claim 44, wherein, in a case that the request is related to a stateful transaction, the processor determines that at least a  
5 second request having a URI substantially the same as the URI of the request should be processed in the network server.

FIG. 1



2/6

FIG.2A



3/6

FROM FIG.2A

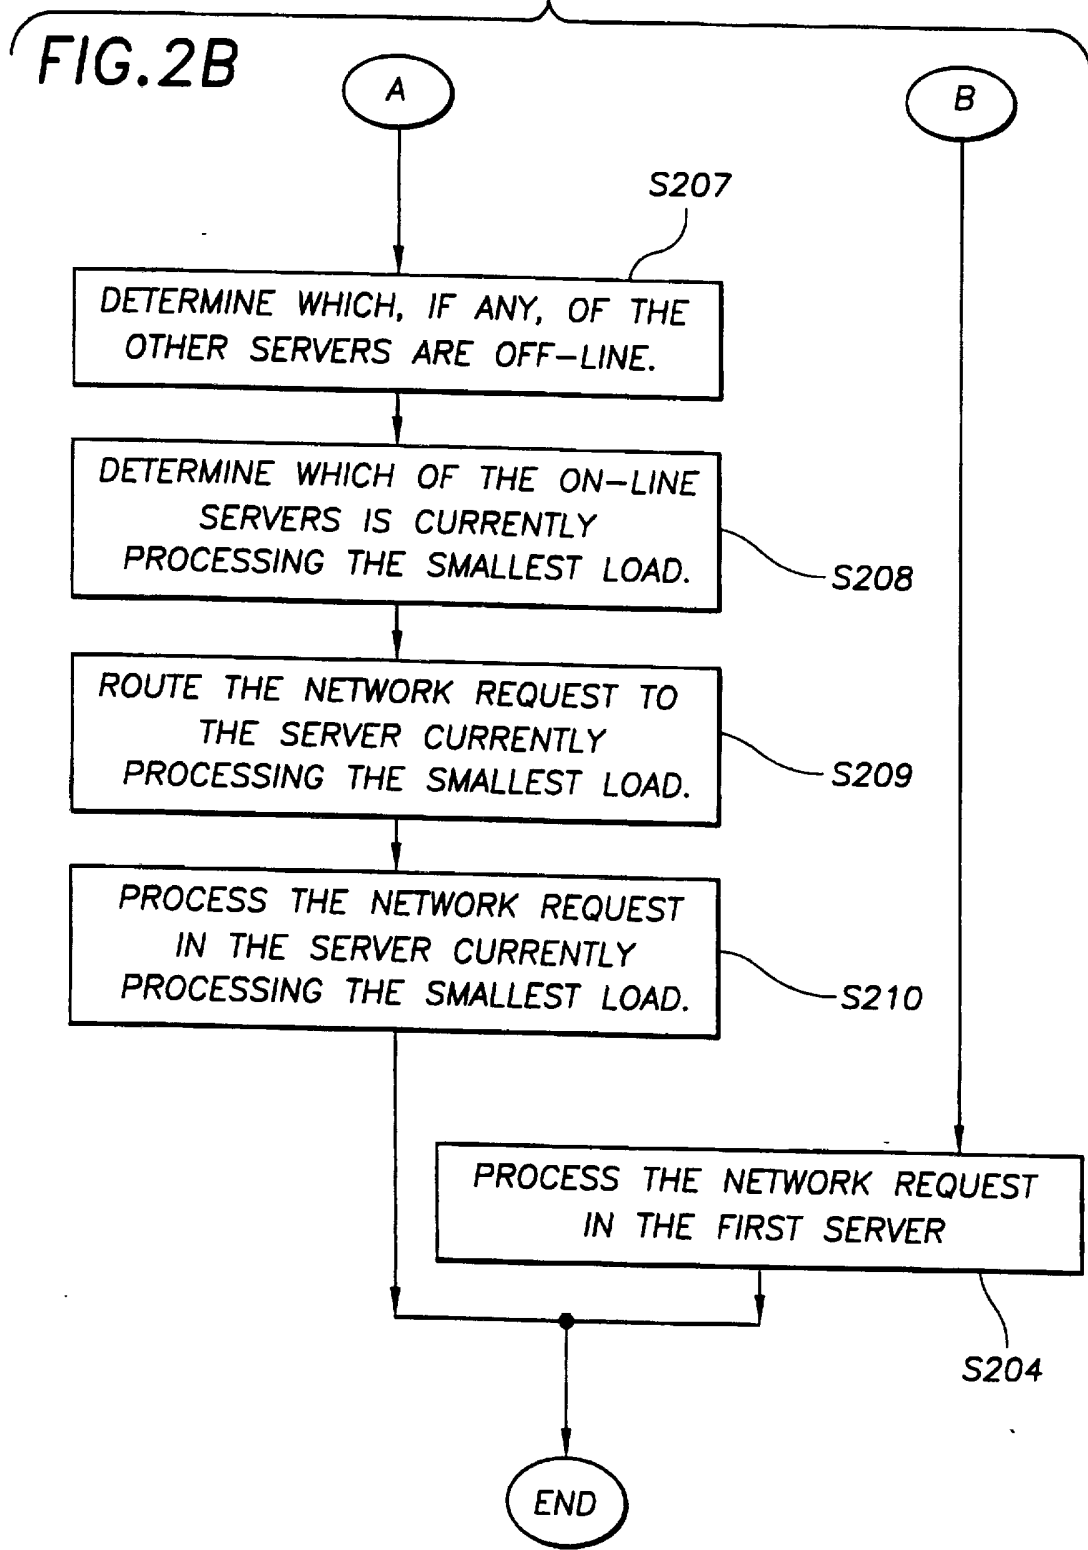


FIG. 3 4/6

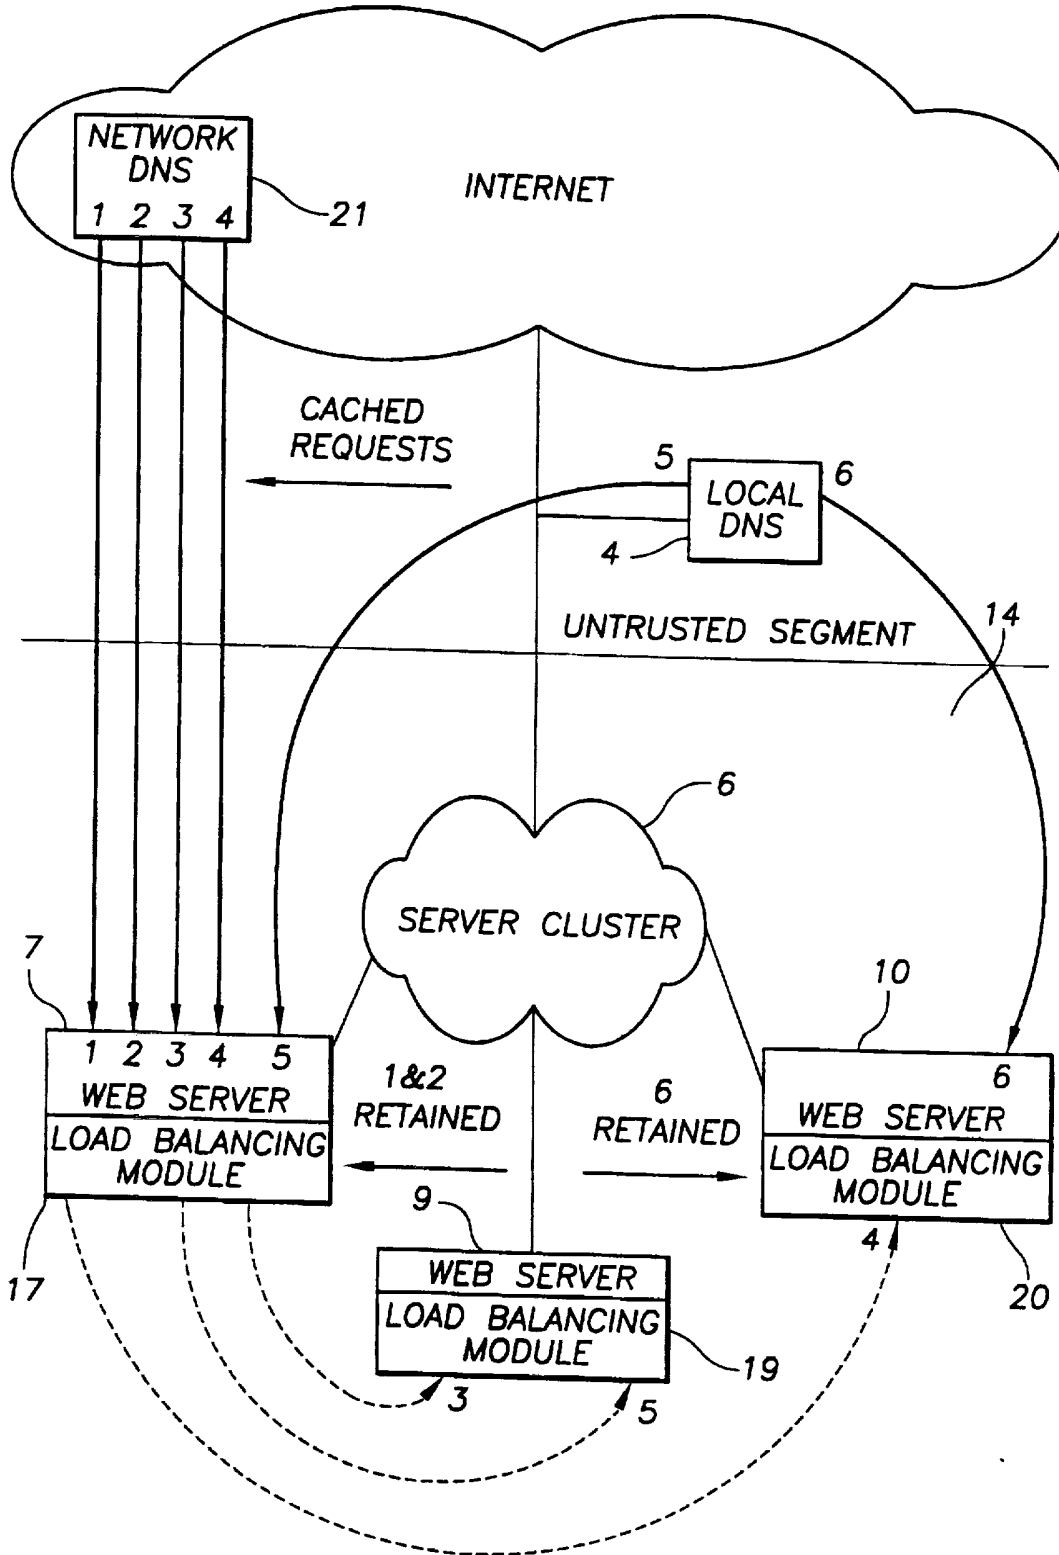




FIG. 4

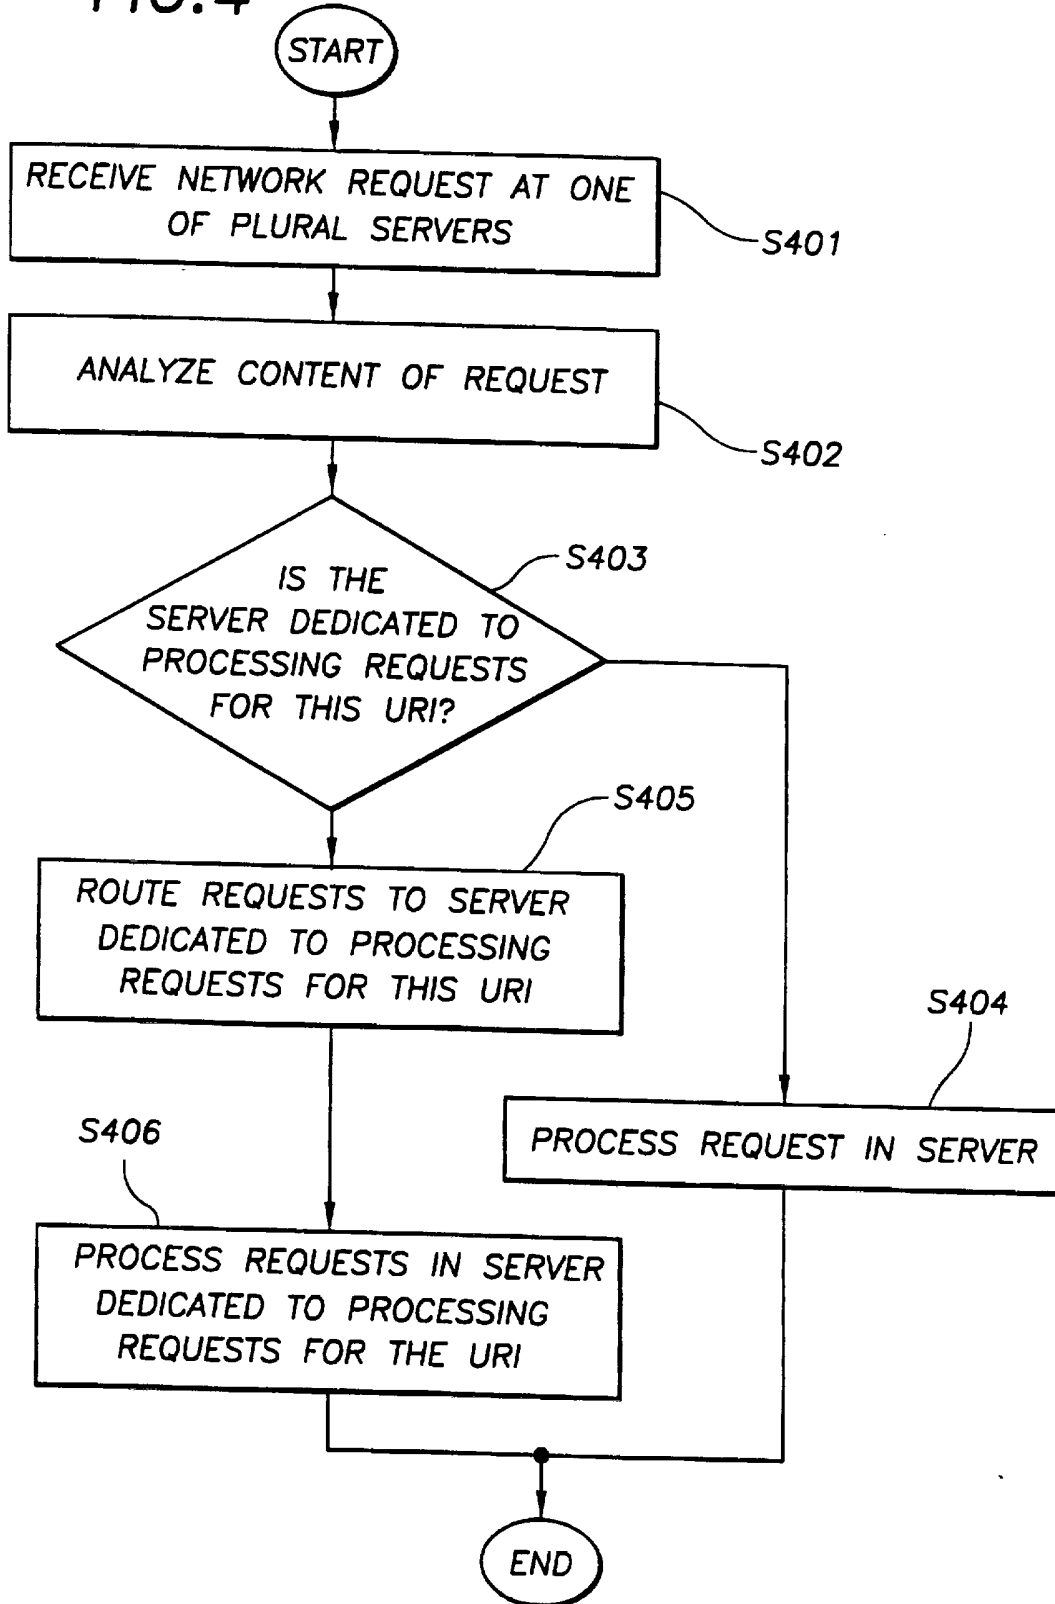


FIG. 5

