

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
2 December 2004 (02.12.2004)

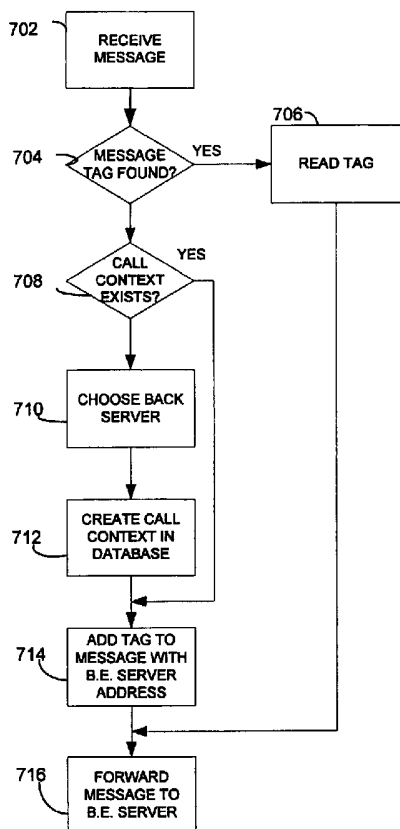
PCT

(10) International Publication Number  
**WO 2004/105341 A1**

- (51) International Patent Classification<sup>7</sup>: **H04L 29/06**, G06F 9/46
- (21) International Application Number: PCT/EP2004/005487
- (22) International Filing Date: 18 May 2004 (18.05.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 03291191.9 20 May 2003 (20.05.2003) EP
- (71) Applicant (for all designated States except US): **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.L.P.** [US/US]; 20555 S.H 249, Houston, TX 77070 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **FORISSIER, Jerome** [FR/FR]; "Le Grand Large A", 55, Av de Cannes, F-06160 Juan Les Pins (FR). **GAYRAUD, Richard** [FR/FR]; Res. Santa Monica, Bat C, 70, av. Maurice Cheval, F-06110 Le Cannet (FR). **LAMBERTON, Marc** [FR/FR]; 981, route de St Jean, F-06600 Antibes (FR). **MANSUTTI, David** [FR/FR]; Les Terrasses de Vallauris/Appt 22 Bt A, 190, route de Cannes, F-06220 Vallauris (FR).
- (74) Common Representative: **LLOYD, Richard**; HP Centre de Competences France, IP Section/ Legal Dept, 5, av. Raymond Chanas, F-38053 Grenoble Cedex 09 (FR).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD,

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR LOAD-BALANCING IN A DISTRIBUTED PROCESSING SYSTEM



(57) Abstract: According to an embodiment of the present invention, there is provided a method of routing a message to one of a plurality of available processing systems, comprising the steps of detecting the presence of a destination identifier in the message, and where the presence of the destination identifier is detected, forwarding the message to the processing system identified thereby. Where the presence of the destination field is not detected the steps of determining a destination processing system for processing the message, inserting into the message a destination identifier identifying the determined destination processing system, and forwarding the message to the determined processing system may be performed.

WO 2004/105341 A1



MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

**Published:**

- with international search report

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

## METHOD AND APPARATUS FOR LOAD-BALANCING IN A DISTRIBUTED PROCESSING SYSTEM

The present invention relates to the field of distributed processing systems and, more particularly, to improvements in load-balancing systems and methods.

5

Increasingly network-based services and applications are accessed via some kind of load-balancing arrangement. For example, when a network-based service has a large number of concurrent users, load-balancing enables the processing load to be distributed among multiple backend servers or applications. Should the number of users increase over time, additional backend servers may be added to cope with the increased load, completely transparently to the user. A load-balancer may, for example, receive all requests to a network-based service, and forward the requests to an appropriate backend server based on the load or some other parameter of the system. One advantage of using load-balancers is that services may be provided with a single externally visible address.

15

Many types of network service, for example those using hypertext transfer protocol (HTTP) over the Internet, use a simple request and response mechanism. For example, an Internet browser sends a HTTP request to an application, and the application responds with a HTTP response. The application address is typically the address of a load-balancer which then routes the request to an available server based, for example, on the current workload of the available servers. In many such situations, all requests may be treated independently of all other requests – including successive requests coming from the same user. In such cases, there is typically no requirement that different requests from the same user be processed by the same backend server. In other words, the requests are effectively contextless and do not require the load-balancer to have any previous knowledge of any previous requests in order to determine to which backend server to forward the message.

20

25

In other situations, context information must be maintained by the load-balancer in order to determine to which backend server a message should be forwarded to. For example, in telecommunication networks, multiple messages may be sent to and from a terminal to establish a call with another party. It is thus generally important that different messages related to the same call be processed by the same server.

30

200301168-EP

For example, when establishing a telephone call it is typically desirable that all messages relating to the call establishment are processed by the same backend server. This type of functionality is commonly referred to as server affinity.

5 In order to provide server affinity a load-balancer is required to store some context information, for example the call ID of each message, and to check each received message to determine whether that call ID is already being processed by a given server, to thereby ensure that subsequent messages having the same call ID are processed by the same backend server. This is typically achieved by maintaining a  
10 database of current calls IDs and backend servers processing messages relating to each call ID. Aspects of traditional telephony systems typically function in this way.

In traditional telephony applications server affinity may only be required during the establishment of a call, which typically only lasts for a few seconds. Thus, once a call  
15 has been established, the load-balancer can remove all references to that call ID, freeing up space in the database.

However many other systems, including some of the new telephony protocols, such as the session initiation protocol (SIP), function in sufficiently different ways that the  
20 conventional load-balancing approach is no longer suitably adapted.

SIP is an application-layer control or signaling protocol for establishing, modifying, and terminating real-time calls and conferences over, primarily, Internet protocol (IP) networks. At its simplest, a call setup in SIP requires two transactions: one to setup a  
25 call; and one to release the call. A SIP transaction typically only lasts a few seconds, whereas a SIP call is theoretically unbounded. Due to the nature of SIP, where, for example, call legs may be added to a current call, media types may be changed at any time, and so on, it is generally required that SIP messages belonging to the same call be processed by the same backend server.

30 Traditional load-balancing approaches, such as maintaining a database of context information for all active calls at a load-balancer, may therefore, not be suitably adapted for use with SIP for a number of reasons. Firstly, since the length of a SIP call is theoretically unbounded and the fact that the load-balancer must store call

context information for the duration of each and every the call, it is possible that the load-balancer may become overwhelmed, especially if the number of simultaneous calls and the number of backend servers is large. Additionally, in order to clean up the stored context information, the load-balancer may be required to receive all incoming and outgoing SIP messages to determine when a call has finished in order to be able to remove the stored context information for that call when it is no longer required. These constraints may impact the performance capabilities of such a load-balancer and limit the message handling capabilities thereof.

10 Accordingly, one aim of the present invention is to provide a system which overcomes at least some of the above-mentioned problems.

According to a first aspect of the present invention, there is provided a method of routing a message to one of a plurality of available processing systems, comprising the steps of detecting the presence of a destination identifier in the message, and where the presence of the destination identifier is detected, forwarding the message to the processing system identified thereby. Where the presence of the destination field is not detected the further steps of determining a destination processing system for processing the message, inserting into the message a destination identifier identifying the determined destination processing system, and forwarding the message to the determined processing system may be performed.

Advantageously, this removes the need for a load-balancing element to maintain a database of all current calls being processed, thus reducing the processing load of the load-balancer. Furthermore, since the load-balancer no longer has to receive all messages sent from backend servers in order to know when to clean up the database, throughput of the load-balancer may be further increased.

Each message may further include a message identifier for identifying related messages in which case the method may further comprise maintaining a database of message identifiers for which no destination identifier was detected along with information indicating to which of the available processing systems each message was forwarded to.

Where a message is received without a destination identifier, the method may further comprise searching the database for a related message identifier and, where found, forwarding the message to the processing system identified therein.

- 5 In one embodiment each processing system is adapted for sending a response to the message originator via the load-balancer. Alternatively each processing system may be adapted for sending a response directly to the originator of the message.

10 Preferably entries are removed from the database after a predetermined amount of time.

The method may be used in conjunction with a session initiation protocol (SIP) based network.

- 15 The destination identifier are preferably inserted into an extension header of a SIP message.

When used in conjunction with a SIP network, the entries may be removed after 32 seconds.

- 20 In one embodiment the method may be adapted for use with the user datagram protocol (UDP).

25 In a further embodiment the method may be adapted for use with a load-balancing element.

- 30 According to a second aspect of the present invention, there is provided a load-balancing system for routing a message to one of a plurality of available processing systems. The system comprises a message analyzer for detecting the presence of a destination identifier in the received message, and a message forwarder for forwarding the message to the processing system identified by the detected identifier.

When the presence of a destination identifier is not detected the system may further comprise a load analyzer for determining a destination processing system for

processing the message, and a message processor for inserting into the message a destination identifier identifying the determined destination processing system.

Each message may further include a message identifier for identifying related  
5 messages, in which case the system further comprises a database for storing details of message identifiers for which no destination identifier was detected along with information indicating to which of the available processing systems each message was forwarded to.

10 Where a message is received without a destination identifier the system further comprises means for searching the database for a related message identifier and for identifying to which processing system the message should be forwarded.

In one embodiment each processing system is adapted for sending a response to the  
15 message originator via the load-balancer. In a further embodiment each processing system is adapted for sending a response directly to the originator of the message.

The load-balancing system may be adapted for use in a session initiation protocol  
(SIP) based network.

20

Preferably the database is adapted to remove entries after a predetermined amount of time. In the case of SIP, the predetermined amount of time is preferably 32 seconds.

25 The message processor preferably inserts the destination identifier into an extension header of a SIP message.

The load-balancing system may be adapted for use with the user datagram protocol  
(UDP).

30

According to a further aspect of the present invention, there is provided a load-balancing element adapted for use with the herein described method.

The invention will now be described, by way of non-limiting example, with reference to the accompanying diagrams, in which:

Figure 1 is a block diagram showing an overview of a simplified SIP network arrangement according to the prior art;

5 Figure 2 is a block diagram showing a back to back user agent (B2BUA) according to the prior art;

Figure 3 is a block diagram showing an overview of a load-balancing system according to an embodiment of the present invention;

Figure 4 is a message flow diagram illustrating a further embodiment;

10 Figure 5 is a message flow diagram illustrating a further embodiment;

Figure 6 is a message flow diagram illustrating a further embodiment;

Figure 7 is a message flow diagram illustrating a further embodiment; and

Figure 8 is a flow diagram outlining example processing steps of a load-balancer operating in accordance with an embodiment of the present invention.

15

Figure 1 is a block diagram showing an overview of a simplified SIP network arrangement 100 according to the prior art. A number of terminals, which may be SIP user agents, 102, 104 and 106 are shown. If a user agent 102 wishes to establish a call with a user agent 104, a SIP Invite message is sent to a SIP proxy server 110 in an IP network 108. The SIP proxy resolves the address of the user agent 104 and forwards the Invite message to a back-to-back user agent (B2BUA) 112. The B2BUA may be used, for example, to provide value added services, such as pre-paid billing or conferencing, in addition to the basic telephony/multimedia functionality provided by SIP. A further SIP proxy 114 is typically used to resolve the address of the user agent 104.

20

25

Typically the value add services mentioned above may be implemented in a distributed processing arrangement, as shown in Figure 2 which shows a B2BUA 112 in accordance with the prior art. A load-balancer 202 is provided which receives all incoming messages sent to the B2BUA 112. As mentioned above, one of the advantages of using a load-balancer is that only a single external network address is seen by the network. The load-balancer 202 forwards each received message to one of a plurality of backend servers 206, 208, 210 which carry out the required processing operations. Such processing operations could include, for example, pre-

30

paid billing, credit card authorizations, portal authentication and so on. The load-balancer 202 may determine which backend server to use through use of any appropriate algorithm, such as least loaded, round-robin and so on. It will be appreciated that although the load-balancer and back-end servers are shown as being integral to the B2BUA 112, these elements may be located externally therefrom, for example in a distributed manner.

Since it is generally advantageous that all SIP messages relating to the same call are processed by the same backend server, the load-balancer 202 has to maintain a database 204 of all current SIP calls along with related call IDs and the backend server which is processing the call. Additionally, all messages sent from a backend server 206, 208 or 210 to a user agent, also pass through the load-balancer 202 so that the load-balancer can determine when a SIP call has terminated thereby allowing the load-balancer to clean up the database 204. As previously mentioned, due to the theoretically endless nature of a SIP call coupled with the large number of simultaneous calls which can be handled by the B2BUA 112, the database 204 has to be sufficiently large to handle data from the maximum number of simultaneous calls which the B2BUA 112 can support.

An embodiment of the invention will now be described with reference to Figure 3 and Figure 8. Figure 3 is a block diagram showing an overview of a load-balancing system according to an embodiment of the present invention and Figure 8 is a flow diagram outlining example processing steps which may be made by the load-balancer 302 of Figure 3. In the following description, for simplicity of explanation, some of the steps shown in Figure 8 are initially not discussed.

The system 300 may be, for example, a B2BUA, or other network or client/server element. A load-balancer 302 receives a message (step 702), for example a SIP INVITE message, from, for example, a SIP user agent 102. When a message arrives at the load-balancer relating to a new call, the load-balancer chooses (step 710) one of the available backend servers 206, 208 or 210 to send the message to using a suitable load-balancing algorithm, as will be appreciated by those skilled in the art. The load-balancer 302 inserts a tag in the SIP message (step 714), for example 'myTag', indicating the identity of the chosen backend server and forwards the

message to the chosen backend server (step 716). Preferably the tag contains sufficient information to enable the load-balancer to route the message without requiring a further call context. Additionally, the tag is preferably inserted into the message such that the tag will be included in all future messages sent in response to the message. For example, in SIP the message may be suitably inserted as an extension header.

SIP provides for the re-transmission of messages in the event that a response is not received within a predetermined amount of time. One problem that this can create is that if a backend server is slow to respond, or if an initial message is lost, for example, the SIP user agent may retransmit the same message. If this message happens to be the first message related to a call (i.e. there is no tag present), the load-balancer is likely to send this message to a different backend server than that dealing with the first message, which may lead to the system creating several call contexts in different backend servers for a single SIP call. This may result in protocol violations for example if identical responses are sent to a user agent client, or sub-optimal processing.

In order to prevent this, the load-balancer 302 preferably maintains a database 304 of all messages which relate to new calls. Thus, the following additional steps are performed. For example, when a message is received it is determined (Step 704) whether a previously inserted tag is present. If not, this indicates that the received message may relate to a new call. The database 304 is searched to determine whether a message with the same call identification is present in the database (step 708). If yes, then the message may be, for example, a retransmitted message or a CANCEL message sent shortly after an initial INVITE message, and is forwarded to the backend server indicated in the database once a suitable tag has been inserted in the message. If no message having the same call ID is found, this indicates that this is the first message relating to a call, in which case a suitable backend server is chosen to process the message (step 710). A call context is subsequently created in the database (step 712), a tag is added to the received message identifying the chosen backend server (step 714), and finally the message is forwarded to the chosen backend server.

After a predetermined amount of time from their creation in the database, 32 seconds in the case of SIP, entries in the databases may be deleted since after this time no further untagged messages relating to the same transaction may be accepted by the load-balancer. Advantageously, for example when user datagram protocol (UDP) is used as the transport protocol, the backend servers may respond directly to the SIP user agents, removing the need to pass through the load-balancer, thus further increasing the potential throughput of the load-balancer. This is possible since the load-balancer maintains no context information after a predetermined time, and therefore does not need to be informed when a call ends in order to clean-up the database.

As will be appreciated by those skilled in the art, the herein-described functionality performed by the load-balancer may be provided in a number of ways, for example, by way of software, by suitable electronic hardware, or a combination of both software and hardware. For example, the load-balancer 302 may comprise suitable logical or functional elements such as a message analyzer, for analyzing the messages to determine whether an inserted tag is present, a load analyzer, for determining to which back-end server a message should be processed by, a message processor, for inserting identification tags into a message, and a message forwarder for forwarding a message to an appropriate backend server. Such elements may be provided in various combinations.

A number of further embodiments will now be described with reference to the message flow diagrams of Figures 4 to 6.

As shown in Figure 4, a SIP user agent (CL1) sends a SIP INVITE message to an application (APP1). As described above, the application is an application using a load-balancer and a number of distributed backend servers. When the load-balancer receives the INVITE message, it analyses the message to determine the presence of an inserted tag. If no tag is found, the load-balancer chooses a suitable backend server (BE1) and inserts a tag, for example *myTag*, identifying the chosen backend server into the message, and forwards the message to the chosen backend server. When the backend server sends a response message to the user agent the message passes first through the load-balancer 302 which forwards the message to the SIP

user agent (CL1). If the SIP user agent sends a further message related to the same call, the previously inserted tag will be present in the message. When this message is received by the load-balancer 302 the message is examined to determine whether there is a tag present therein. If there is a tag present, the load-balancer forwards the message to the backend server indicated by the tag. In this way, the load-balancer does not need to store details of all current calls in order to determine which backend server should process a message and the requirement to store context information for all calls at the load-balancer is removed.

10 In this embodiment, no modifications are required to either the SIP user client or the backend servers, as the effect of inserting the tag is effectively transparent. Furthermore, in the event of a failure of a backend server, the load-balancer can forward the message to a backup server other than that indicated by the inserted tag, without the SIP user agent ever being aware that a failure occurred.

15

The message flow diagram of Figure 5 illustrates the way in which the retransmission of a SIP INVITE message may be handled by the load-balancer. A SIP user agent (CL1) sends a SIP INVITE message to the load-balancer (APP1). If the load-balancer determines that no inserted tag is present in the received message, for example if the message is the first message relating to a particular call identification, the load-balancer first searches its database to see if any entries have a matching call identification. If an entry is found this may indicate, for example, that the received message is a retransmitted message. The load-balancer retrieves the address of the backend server from the database, inserts a corresponding tag into the message, and forwards the message to the appropriate backend server.

20

If no matching entries are found this implies that the received message is the first message of a call and the message is processed as described above.

30 As described previously, a predetermined amount of time after the details of the first SIP message are stored in the database all entries having the corresponding call identification may be erased. In this way, the database only contains context information for a given call identification for a maximum of 32 seconds.

A still further embodiment is illustrated with reference to Figure 6. In this embodiment, the backend server is adapted to respond directly to the SIP user agent. This is possible, for example, when the transport protocol used is user datagram protocol (UDP). Preferably the SIP message sent from the backend server is such that a  
5 response message sent from the user agent will be addressed to the load-balancer. This is particular useful for fault tolerance purposes as the load-balancer maintains full control of the routing of messages, for example, in the event of a backend server failing.

10 In some circumstances, a backend server may initiate a call with a SIP use agent, as illustrated with reference to Figure 7. In this embodiment, before sending the invite message, a tag is inserted in the message for enabling the load-balancer to later identify the backend server which initiated the call. Thus, when the SIP user agent sends a response message via the load-balancer, the load-balancer is able to  
15 determine which backend server to forward the message without requiring any context information to be stored at the load-balancer.

In such a system as described above the resource requirements of a load-balancer are no longer proportional to the number of established calls, since only the context  
20 information of newly established calls is required to be maintained by the load-balancer.

**CLAIMS**

1. A method of routing a message to one of a plurality of available processing  
5 systems, comprising the steps of:  
    detecting the presence of a destination identifier in the message;  
    and  
    where the presence of the destination identifier is detected, forwarding the  
message to the processing system identified thereby;  
10 and where the presence of the destination field is not detected:  
    determining a destination processing system for processing the message;  
    inserting into the message a destination identifier identifying the determined  
destination processing system; and  
    forwarding the message to the determined processing system.
- 15
2. A method according to claim 1, wherein each message further includes a message  
identifier for identifying related messages, the method further comprising maintaining  
a database of message identifiers for which no destination identifier was detected  
along with information indicating to which of the available processing systems each  
20 message was forwarded to.
3. The method according to claim 2, further comprising, where a message is received  
without a destination identifier, searching the database for a related message  
identifier and, where found, forwarding the message to the processing system  
25 identified therein.
4. The method according to claim any previous claim, wherein each processing  
system is adapted for sending a response to the message originator via the load-  
balancer.
- 30
5. The method according to any of claims 1 to 4, wherein each processing system is  
adapted for sending a response directly to the originator of the message.

6. The method according to any previous claim, further comprising removing entries in the database after a predetermined amount of time.
7. The method according to any previous claim, for use in a session initiation protocol (SIP) based network.
8. The method according to claim 7, wherein the step of inserting further comprises inserting the destination identifier into an extension header of a SIP message.
9. The method according to claim 7 or 8, wherein the step of removing entries is adapted to remove entries after 32 seconds.
10. The method according to any previous claim, adapted for use with the user datagram protocol (UDP).
11. The method according to any previous claim, adapted for use with a load-balancing element.
12. A load-balancing system for routing a message to one of a plurality of available processing systems, comprising:  
a message analyzer for detecting the presence of a destination identifier in the received message;  
and a message forwarder for forwarding the message to the processing system identified by the detected identifier.
13. A load-balancing system according to claim 12, further comprising, for when the presence of a destination identifier is not detected,  
a load analyzer for determining a destination processing system for processing the message; and  
a message processor for inserting into the message a destination identifier identifying the determined destination processing system.
14. A load-balancing system according to claim 12 or 13, wherein each message further includes a message identifier for identifying related messages, and further

comprising a database for storing details of message identifiers for which no destination identifier was detected along with information indicating to which of the available processing systems each message was forwarded to.

5 15. A load-balancing system according to claim 14, further comprising, where a message is received without a destination identifier, means for searching the database for a related message identifier and for identifying to which processing system the message should be forwarded.

10 16. A load-balancing system according to any of claims 12 to 15, wherein each processing system is adapted for sending a response to the message originator via the load-balancer.

15 17. A load-balancing system according to any of claims 12 to 16, wherein each processing system is adapted for sending a response directly to the originator of the message.

18. A load-balancing system according to any of claims 12 to 17, wherein the database is adapted to remove entries after a predetermined amount of time.

20

19. A load-balancing system according to any of claims 12 to 18, adapted for use in a session initiation protocol (SIP) based network.

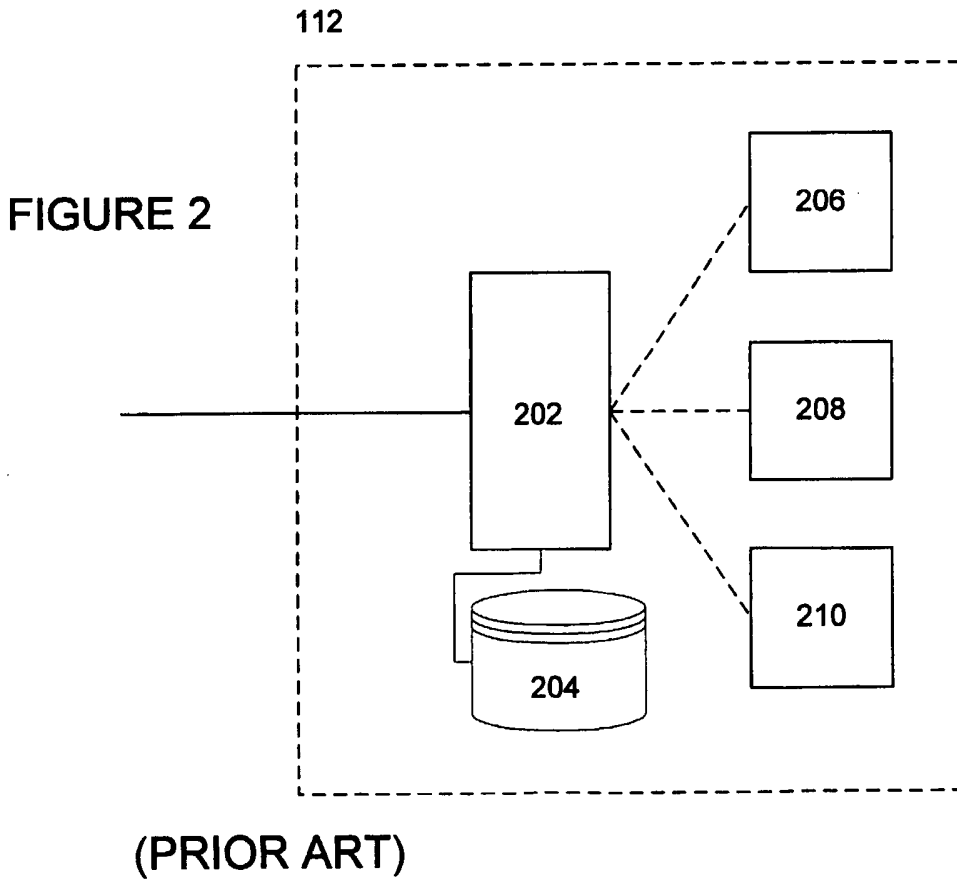

25 20. A load-balancing system according to claim 19, wherein the message processor is adapted to insert the destination identifier into an extension header of a SIP message.


21. A load-balancing system according to claim 19 or 20, wherein the predetermined length of time is 32 seconds.

30

22. A load-balancing system according to any of claims 12 to 21, for use with the user datagram protocol (UDP).

23. A load-balancing element adapted for use in accordance with any of claims 1 to 11.





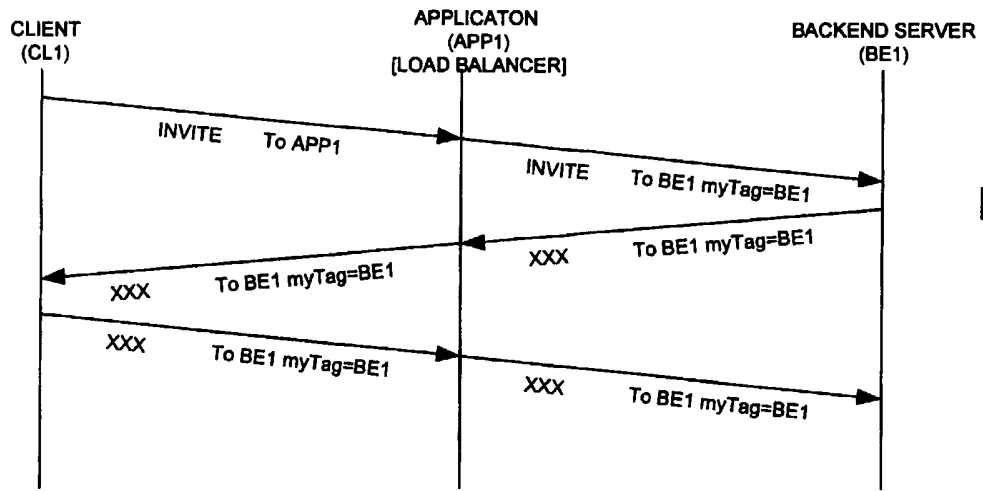


FIG 4

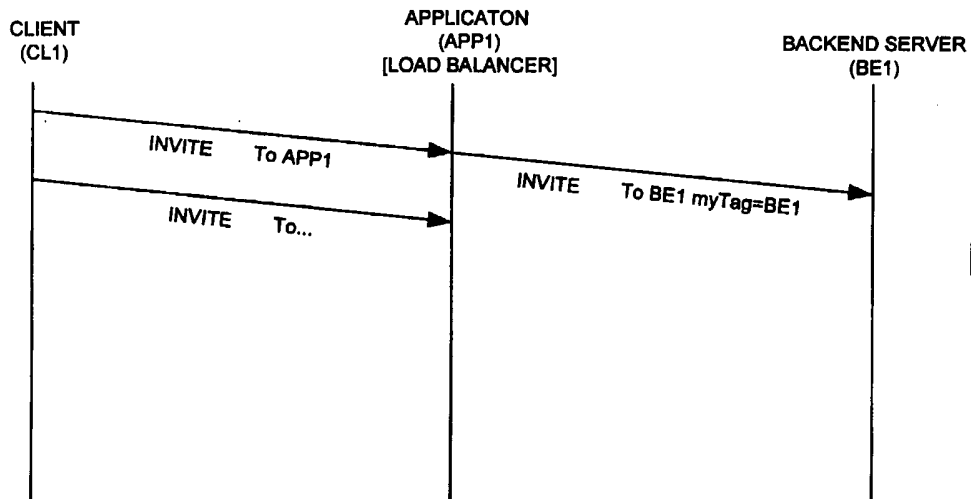


FIG 5

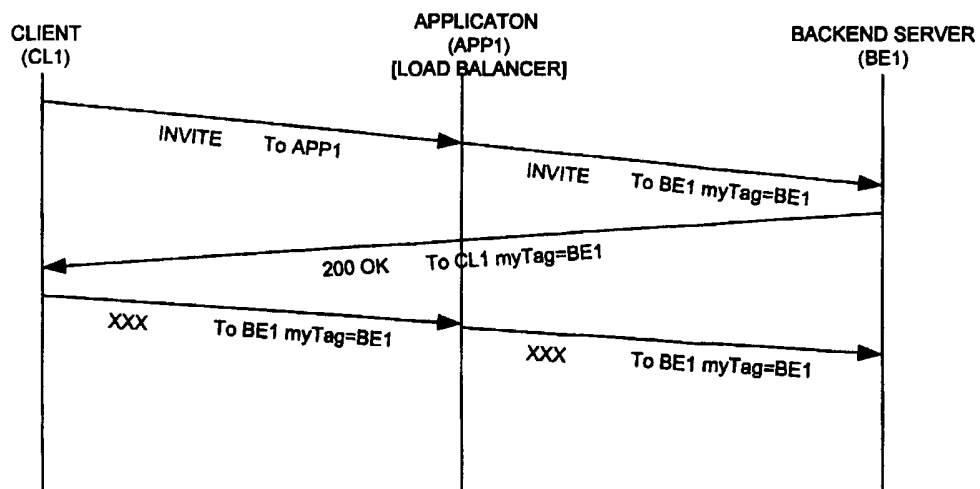


FIG 6

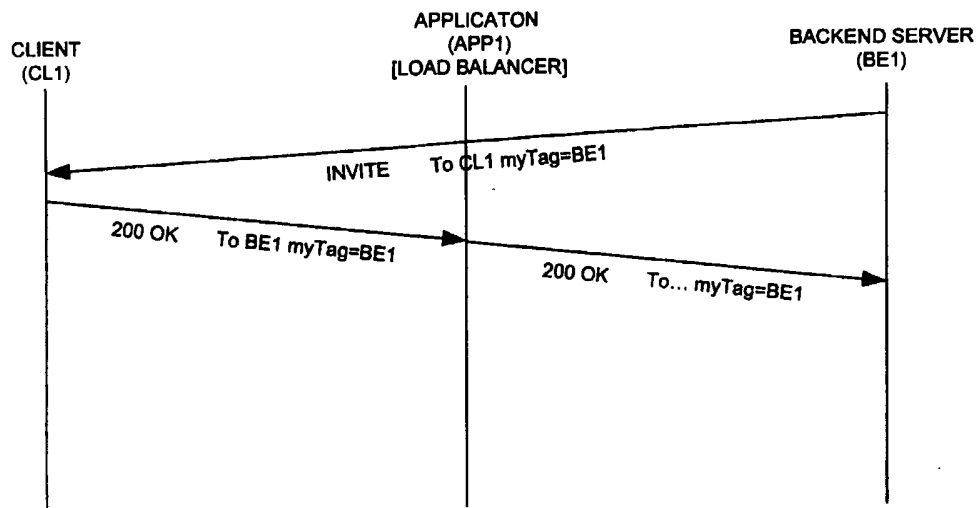
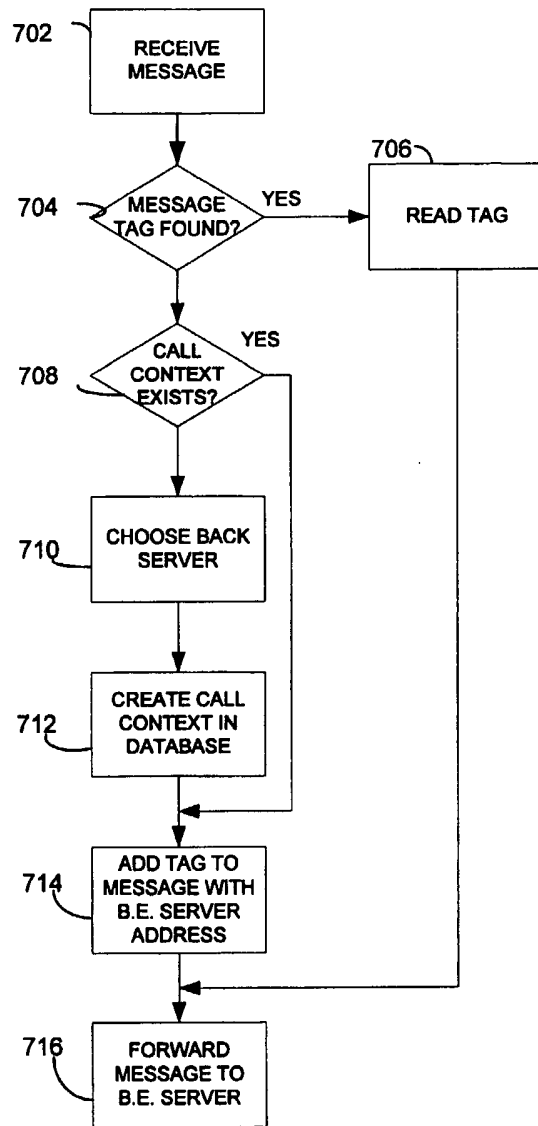


FIG 7

FIGURE 8



# INTERNATIONAL SEARCH REPORT

International Application No  
PCT/EP2004/005487

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 7 H04L29/06 G06F9/46

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

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
IPC 7 H04L G06F

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

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

EPO-Internal, WPI Data, INSPEC

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 924 097 A (COOPER THOMAS PEARSON ET AL) 13 July 1999 (1999-07-13) column 7, line 8 - column 8, line 63 -----	1-23
A	US 6 345 303 B1 (TSO MICHAEL MAN-HAK ET AL) 5 February 2002 (2002-02-05) column 6, line 62 - column 8, line 53 -----	1-11, 13-23
X		12
A	US 2003/093462 A1 (VAINIKAINEN MIKKO ET AL) 15 May 2003 (2003-05-15) paragraph '0034! - paragraph '0046! -----	1,7,8, 12,19,20

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

\* Special categories of cited documents:

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

- \*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \*&\* document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

10 August 2004

17/08/2004

Name and mailing address of the ISA  
European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer  
  
Jurca, A

# INTERNATIONAL SEARCH REPORT

International Application No  
PCT/EP2004/005487

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5924097	A	13-07-1999	NONE
US 6345303	B1	05-02-2002	NONE
US 2003093462	A1	15-05-2003	WO 03042852 A1 22-05-2003