

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 June 2007 (21.06.2007)

PCT

(10) International Publication Number
WO 2007/068664 A1

(51) International Patent Classification:
H04L 29/08 (2006.01)

(21) International Application Number:
PCT/EP2006/069496

(22) International Filing Date:
8 December 2006 (08.12.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/164,932 12 December 2005 (12.12.2005) US

(71) Applicant (for all designated States except US): INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(71) Applicant (for MG only): IBM UNITED KINGDOM LIMITED [GB/GB]; PO Box 41, Portsmouth Hampshire PO6 3AU (GB).

(72) Inventor; and

(75) Inventor/Applicant (for US only): NUSBICKEL, Wendi, Lynn [US/US]; 824 Dover Street, Boca Raton, Florida 33487 (US).

(74) Agent: WALDNER, Philip; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).

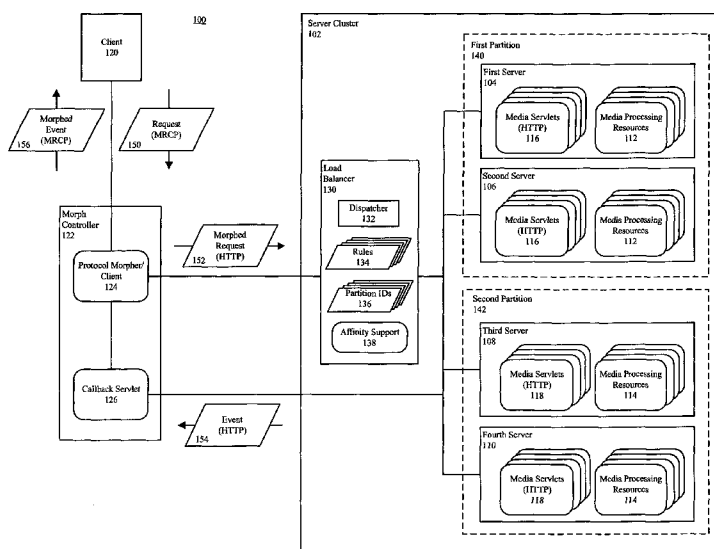
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR,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

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

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.

(54) Title: LOAD BALANCING AND FAILOVER OF DISTRIBUTED MEDIA RESOURCES IN A MEDIA SERVER



(57) Abstract: A method and a system for load balancing media resources among a plurality of servers. A first request can be received from a client requesting a first media processing resource, the request being formatted in accordance with a first protocol. The first request can be morphed into a first morphed request formatted in accordance with a second protocol. A first server that is least loaded can be selected from a first server partition including a plurality of servers, each of which provides the requested first media processing resource. The first morphed request can be forwarded to the first server in accordance with the second protocol.

WO 2007/068664 A1

**LOAD BALANCING AND FAILOVER OF DISTRIBUTED MEDIA RESOURCES
IN A MEDIA SERVER**

DESCRIPTION

5

Reservation of Rights in Copyrighted Material

10

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

15

Field of the Invention

The present invention relates to distributed resource management and, more particularly, to management of distributed media resources.

20

Background of the Invention

25

Automatic speech recognizers (ASRs) and text to speech (TTS) synthesizers are speech resources commonly implemented in automated interactive systems, for example in automated voice response systems used in telephone call centers. ASRs and TTS synthesizers are resource consumptive, and oftentimes must be implemented across several machines that constitute a speech server solution. In order to reduce hardware costs, the speech resources are typically shared among many clients and allocated only when they are actually needed.

30

35

Media processing resource Control Protocol (MRCP) is a protocol that provides a mechanism for a client device requiring media processing resources, for example ASRs and TTS synthesizers, to control such resources on a network. In particular, MRCP defines requests, responses and events needed to control the media processing resources. Unfortunately, the MRCP standard exclusively relies on the client to distribute client requests. Indeed, MRCP lacks methods that can be implemented by a server to distribute the media processing resources across multiple nodes. Moreover, clients accessing resources may not have

knowledge of which server resources are available and which server
resources are currently being used. Accordingly, the most common form of
load balancing is a method in which a client sequentially sends requests
to various servers until a client session is established. Such a method,
5 however, does not properly balance server loads, which reduces hardware
efficiency.

It therefore would be beneficial to provide a technique which enables an
MRCP server cluster to distribute its own media processing resources and
10 to load balance such resources.

Summary of the Invention

The present invention relates to a method and a system for load balancing
15 media resources among a plurality of servers. A first request can be
received from a client requesting a first media processing resource, the
request being formatted in accordance with a first protocol. The first
request can be morphed into a first morphed request formatted in
accordance with a second protocol. A first server that is least loaded
20 can be selected from a first server partition including a plurality of
servers, each of which provides the requested first media processing
resource. The first morphed request can be forwarded to the first server
in accordance with the second protocol.

25 Another embodiment of the present invention can include a machine readable
storage being programmed to cause a machine to perform the various steps
described herein.

Brief Description of the Drawings

30 Preferred embodiments of the present invention will be described below in
more detail, with reference to the accompanying drawings, in which:

FIG. 1 is a block diagram illustrating a system in which media processing
35 resources are load balanced in accordance with one embodiment of the
present invention.

FIG. 2 is a flow chart illustrating a method of load balancing media processing resources in accordance with another embodiment of the present invention.

5 **Detailed Description**

While the specification concludes with claims defining the features of the invention that are regarded as novel, it is believed that the invention will be better understood from a consideration of the description in
10 conjunction with the drawings. As required, detailed embodiments of the present invention are disclosed herein; however, it is to be understood that the disclosed embodiments are merely exemplary of the invention, which can be embodied in various forms. Therefore, specific structural and functional details disclosed herein are not to be interpreted as
15 limiting, but merely as a basis for the claims and as a representative basis for teaching one skilled in the art to variously employ the present invention in virtually any appropriately detailed structure. Further, the terms and phrases used herein are not intended to be limiting but rather to provide an understandable description of the invention.

20 The embodiments disclosed herein relate to a method and a system which may be used for load balancing Media Resource Control Protocol (MRCP) resources among a plurality of servers. For example, an MRCP message requesting a media processing resource can be received over a TCP/IP
25 socket by a protocol morpher/client, and morphed into a standard Hypertext Transfer Protocol (HTTP) message. The HTTP message then can be forwarded to a load balancer. From a set of servers, each of which provides the requested media processing resource, the load balancer can select a server that is least loaded and forward the morphed request to the selected
30 server in accordance with the HTTP protocol. When a next request is received by the load balancer, the load balancer can again select a least loaded server from the set of servers and send the next request to that server. The server selected the second time may or may not be the same server that was previously selected.

35 FIG. 1 is a block diagram illustrating a system 100 that includes a server cluster 102. The server cluster 102 can include a plurality of servers, for example a first server 104, a second server 106, a third server 108 and a fourth server 110. Each of the servers 104, 106, 108, 110 can be,

for instance, protocol based servers, such as MRCP servers. The servers 104, 106, 108, 110 can be networked in any suitable manner. For example, the servers can be networked via a local area network (LAN), a wide area network (WAN), the Internet, or any other communications network.

5

Each of the servers 104, 106, 108, 110 can provide media processing resources. For instance, the first server 104 and the second server 106 each can provide a first set of media processing resources 112.

10

Similarly, the third server 108 and the fourth server 110 can provide a second set of processing resources 114. The media processing resources 112, 114 can be automatic speech recognizers (ASRs), text to speech (TTS) synthesizers, speaker verifiers, dual tone multi-frequency (DTMF) recognizers, or any other resources that can process media requests. For example, the first set of media processing resources 112 can include text to speech (TTS) synthesizers that generate a female voice, and the second set of media processing resources 114 can include TTS synthesizers that produce a male voice.

15

20

The system 100 also can include a plurality of clients, such as client 120. The client 120 can be any device that can connect to the server cluster 102 over a network and to which the server cluster 102 provides media processing resources. For example, the client 112 can be a computer, a telephone, a mobile telephone, a personal digital assistant, a game console, an interactive appliance, or any other device which can access media processing resources provided by the server cluster 102 over a LAN, a WAN, the Internet, or any other communications network. Although one client 120 is shown, the system 100 can include any number of clients accessing the media processing resources 112, 114. Indeed, the methods described herein for distributing media processing resources 112, 114 facilitate simultaneous use of the media processing resources 112, 114 by multiple clients.

25

30

35

The client 120 can communicate with the server cluster 102 by propagating a first message, such as a request 150, to a morph controller 122 communicatively linked to the server cluster 102. The request 150 can be formatted in accordance with a suitable communications protocol. In addition, the client 120 can receive a second message, such as an event 156, in response to the first message. In one arrangement, the messages

can be formatted, and communicated to between the client 120 and the morph controller 122, in accordance with MRCP.

To communicate messages in accordance with MRCP version 1, the client 120 can establish a user session with the morph controller 122 using a
5 protocol such as the Real Time Streaming Protocol (RTSP), which operates over TCP/IP. Communicating messages in this manner is generally referred to as "tunneling". To communicate messages in accordance with MRCP
10 version 2, the client 120 can establish a user session using the Session Initiation Protocol (SIP). In contrast to RTSP, after the SIP user session is established, messages are sent directly over TCP/IP instead of being tunneled.

The morph controller 122 can be realized in hardware, software, or a combination of hardware and software. The morph controller 122 also can
15 be realized in a centralized fashion in one processing system or in a distributed fashion where different elements are spread across several interconnected processing systems. Although shown in FIG. 1 as a component distinct from the server cluster 102, the morph controller 122 also can be included as a component of the server cluster 102.

The morph controller 122 can comprise a protocol morpher/client 124 for receiving the request 150 and morphing the request into a morphed request 152. The morphed request 152 can be formatted in accordance with a
25 protocol that is different than the protocol with which the original request 150 is formatted. For example, the protocol morpher/client 124 can morph the request 150 from an MRCP request into an HTTP request. To morph the request 150, MRCP specific code in the request can be replaced with correlating HTTP specific code. Table 1 below illustrates an example of code that can be contained in the request 150 prior to being morphed.

30 **TABLE 1** - Request Example

```
ANNOUNCE rtsp://localhost/media/synthesizer RTSP/1.0  
CSeq: 1  
35 Session 1.IBM.9.22.74.38  
Date: Tue, 15 FEB 2005 11:28:01 est  
Content-Type: application/mrcp  
Content-Length: 149
```

SPEAK 100 MRCP/1.0
 voice-name: Andrew
 Content-Type: application/synthesis+ssml
 Content-Length: 45

5

```
<?xml version="1.0"?>
<say>1 2 3 .</say>
```

10

In order to morph the request illustrated in Table 1, the protocol morpher/client 124 can change code in the RTSP portion of the message as follows:

replace

```
"ANNOUNCE rtsp://localhost/media/synthesizer RTSP/1.0"
```

with

15

```
"POST /synthesizerServlet HTTP 1.1"
"Host: myserver.bocaraton.ibm.com"
"user-agent=Java/1.4.1"
"accept=text/html, image/gif, image/jpeg, *; q=.2, */*; q=.2"
"connection=keep-alive".
```

20

In this arrangement, the MRCP portion of the message body can remain unchanged. Table 2 below illustrates an example of the morphed request 152 generated by morphing the first request 150 in this manner.

25

TABLE 2 - Morphed Request Example #1

```
POST /synthesizerServlet HTTP 1.1
Host: myserver.bocaraton.ibm.com
user-agent=Java/1.4.1
accept=text/html, image/gif, image/jpeg, *; q=.2, */*; q=.2
connection=keep-alive
CSeq: 1
Session 1.IBM.9.22.74.38
Date: Tue, 15 FEB 2005 11:28:01 est
Content-Type: application/mrcp
Content-Length: 149
```

35

SPEAK 100 MRCP/1.0
 voice-name: Andrew

Content-Type: application/synthesis+ssml
 Content-Length: 45

<?xml version="1.0"?>
 5 <speake>1 2 3 .</speake>

In another arrangement the protocol morpher/client 124 can change code in the RTSP portion of the message as follows:

replace

10 "ANNOUNCE rtsp://localhost/media/synthesizer RTSP/1.0"

with

"POST /synthesizerServlet HTTP 1.1"

"Host: myserver.bocaraton.ibm.com"

"user-agent=Java/1.4.1"

15 "accept=text/html, image/gif, image/jpeg, *; q=.2, */*; q=.2"

"connection=keep-alive"

and remove

"Content-Type: application/mrcp"

"Content-Length: 149".

20 In this arrangement code in the MCRP portion of the message can be changed as follows:

replace

"SPEAK 100 MRCP/1.0"

with

25 "MRCPMethod: Speak"

Morphing the code in this manner combines the RTSP and MRCP messages, which changes "SPEAK 100" to an "MRCPMethod: Speak". In addition, the functionality of the content-type and content-length headers that are removed from the RTSP portion of the message can be handled by similar code already contained in the MRCP portion of the message. Table 3 below illustrates an example of the morphed request 152 generated by morphing the request 150 in accordance with this example.

TABLE 3 - Request Example #2

35 POST /synthesizerServlet HTTP 1.1
 Host: myserver.bocaraton.ibm.com
 user-agent=Java/1.4.1
 accept=text/html, image/gif, image/jpeg, *; q=.2, */*; q=.2

connection=keep-alive
CSeq: 1
Session 1.IBM.9.22.74.38
Date: Tue, 15 FEB 2005 11:28:01 est

5

MRCPMethod: Speak
voice-name: Andrew
Content-Type: application/synthesis+ssml
Content-Length: 45

10

```
<?xml version="1.0"?>  
<speake>1 2 3 .</speake>
```

15

By reversing the process described above, the protocol morpher/client 124 also can morph an event 154 generated in response to the morphed request 152 into a morphed event 156 that is formatted in accordance with MRCP. For example, HTTP specific code in the first event 154 can be replaced with correlating MRCP specific code.

20

The protocol morpher/client 124 also can act as a client to the servers 104, 106, 108, 110. For example, the protocol morpher/client 124 can act as an HTTP client. In particular, the protocol morpher/client 124 can propagate the morphed request 152 to a load balancer 130. The load balancer 130 can include a dispatcher 132 that routes requests to the servers 104, 106, 108, 110. Specifically, the dispatcher 132 can route the morphed request 152 to a one of the servers 104, 106, 108, 110 which has media processing resources 112, 114 suitable for responding to the morphed request 152 and which is least loaded among the group of servers having such media processing resources 112, 114.

30

35

For example, the dispatcher 132 can apply resource selection rules 134 to the morphed request 152 to identify a partition of servers, such as a first partition 140 or a second partition 142, that are suitable for responding to the morphed request 152. For instance, the dispatcher 132 can apply the resource selection rules 134 to a pattern in the uniform resource identifier (URI) of the morphed request 152 and/or to a content-type header to identify suitable media processing resources 112, 114. The dispatcher 132 then can reference partition identifiers 136 to determine

which partition, e.g. the first partition 140 or the second partition 142, comprises servers that provide the identified resources.

In a more specific example, the morphed request 152 can include a content-type header that contains "voice=Lisa" and "content type = TTS".

Alternatively, an identifier that indicates "voice=Lisa" and "content type = TTS" can be contained in the URI of the morphed request 152. In response to receiving the morphed request 152, the dispatcher 132 can apply the resource selection rules 134 to determine that the type of media processing resource that provides the "Lisa" TTS is the media processing resource 112. The dispatcher 132 can reference the partition identifiers 136 to select the first partition 140 that includes the first server 104 and the second server 106. The dispatcher 132 then can reference a routing list containing routing information corresponding to each of the servers 104, 106, and query each of the servers 104, 106 to determine which server is least loaded.

If a particular server is no longer available, for instance the server has been taken off-line, that server can be removed from routing list.

Accordingly, the dispatcher 152 can essentially ignore the off-line server. When the server comes back on -line, the server can again be added to the routing list. In this manner, the load balancer 130 also provides failover capability in addition to load balancing.

After determining which server is least loaded, the dispatcher 132 can forward the morphed request 152 to that server. For instance, if the first server 104 is least loaded, the dispatcher 132 can select the URI of the media processing resource 112 on the first server 104. The dispatcher 132 then can call the media servlet 116 on the first server 104 by sending the morphed request 152 to the media servlet 116. The media servlets 114, 116 can be HTTP servlets that are implemented with a suitable middleware platform that supports instantiating the servlets in a distributed fashion across a plurality of server nodes. One example of such a platform is J2EE, but the invention is not limited in this regard.

To achieve affinity with the selected media processing resource 104, the protocol morpher/client 120 can support the client session with cookies and/or uniform resource locator (URL) rewriting. The load balancer 130 also can include affinity support 138 to maintain affinity between the

protocol morpher/client 120 and the media processing resource 104 while required. For example, the load balancer 130 can be configured to passively allow cookies to bypass load balancing routing for desired circumstances, for instance when modifying an outstanding request to a media processing resource. The affinity support also can be configured to obtain affinity to servers for the protocol morpher/client 120, where needed, by adding cookie affinity.

The media servlet 116 can parse the morphed request 152 and forward the parsed request 152 to the selected media processing resource 112. Information parsed from the morphed request 152 can include a URI of a callback servlet 126 to indicate that an event 154 generated by the media processing resource 112 in response to the morphed request 152 is to be forwarded to the callback servlet 126. The parsed information also can include a session identifier to indicate to which client 120 the event 154 is to be forwarded.

The callback servlet 126 can be communicatively linked to the protocol morpher/client 124. The callback servlet 126 can parse the session identifier from event 154 and forward the event 154 to the protocol morpher/client 124, indicating the appropriate client 120 to which the response 156 is to be forwarded. For example, the callback servlet 126 can provide the URL of the client 120. The protocol morpher/client 124 then can morph the event 154 into the morphed event 156 and forward the morphed event 156 to the client 120 in accordance with the first protocol. For example, the protocol morpher/client 124 can forward the morphed event 156 in accordance with MRCP.

When a second request is received from the client 120, or any other client, the process described above for the original request 150 can be repeated. If the second request requires the same media processing resources 112 as the first request, and the first server 104 is still least loaded, the morphed request can again be forwarded to the first server 104. If, however, the first server 104 is no longer the least loaded server in the first partition 140, another server that is least loaded can be selected. For instance, if the second server 106 now is the least loaded server in the first partition 140, the second server 106 can be selected and the second morphed request can be forwarded to it.

If the second request requires different media processing resources than the first request, such as media processing resources 114, the second morphed request can be forwarded to another server having the requested media processing resources 114, for instance the third server 108. The
5 third server 108 can be selected from the second partition 142 over the fourth server 110 if the third server 108 is least loaded of the servers comprising the second partition 142.

FIG. 2 is a flow chart illustrating a method 200 of load balancing media
10 processing resources in accordance with another embodiment of the present invention. Beginning at step 205 a first request can be received from a client requesting a first media processing resource, the request being formatted in accordance with a first protocol. At step 210, the first request can be morphed into a first morphed request formatted in
15 accordance with a second protocol. At step 215, a first server that is least loaded can be selected from a first server partition including a plurality of servers, each of which provides the requested first media processing resource. Proceeding to step 220, the first morphed request can be forwarded to the first server in accordance with the second
20 protocol.

At step 225, a second request can be received from the client. The second request can be morphed into a second morphed request that is formatted in accordance with the second protocol, as shown in step 230. Referring to
25 decision box 235, if the first media processing resource is requested by the second request, the method can continue to decision box 240. If the first server is still the least loaded server in the first partition, the second morphed request can be forwarded to the first server in accordance with the second protocol, as shown in step 245. If, however, the first
30 server is not still the least loaded server, a second server that is least loaded can be selected from the second partition, as shown in step 250. Proceeding to step 260 the second morphed request can be sent to the second server in accordance with the second protocol.

Referring again to decision box 235 and to step 255, if the first media
35 processing resource is not selected by the second request, a second server that is least loaded can be selected from a second server partition comprising a plurality of servers, each of which provides a requested second multimedia processing resource. Referring again to step 260, the

second morphed request can be forwarded to the second server in accordance with the second protocol.

5 The present invention can be realized in hardware, software, or a combination of hardware and software. The present invention can be realized in a centralized fashion in one processing system or in a distributed fashion where different elements are spread across several interconnected processing systems. Any kind of processing system or other apparatus adapted for carrying out the methods described herein is suited.
10 A typical combination of hardware and software can be a general-purpose processing system with an application that, when being loaded and executed, controls the processing system such that it carries out the methods described herein. The present invention also can be embedded in an application product, which comprises all the features enabling the
15 implementation of the methods described herein, and which when loaded in a processing system is able to carry out these methods.

The terms "computer program", "software", "application", variants and/or combinations thereof, in the present context, mean any expression, in any
20 language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following: a) conversion to another language, code or notation; b) reproduction in a different material form. For example, an application can include, but is
25 not limited to, a subroutine, a function, a procedure, an object method, an object implementation, an executable application, an applet, a servlet, a source code, an object code, a shared library/dynamic load library and/or other sequence of instructions designed for execution on a processing system.

30 The terms "a" and "an," as used herein, are defined as one or more than one. The term "plurality", as used herein, is defined as two or more than two. The term "another", as used herein, is defined as at least a second or more. The terms "including" and/or "having", as used herein, are
35 defined as comprising (i.e., open language). The term "coupled", as used herein, is defined as connected, although not necessarily directly, and not necessarily mechanically, i.e. communicatively linked through a communication channel or pathway.

This invention can be embodied in other forms without departing from the spirit or essential attributes thereof. Accordingly, reference should be made to the following claims, rather than to the foregoing specification, as indicating the scope of the invention.

CLAIMS

5 1. A method of load balancing media resources among a plurality of servers comprising:

receiving a first request from a client requesting a first media processing resource, the request being formatted in accordance with a first protocol;

10 morphing the first request into a first morphed request formatted in accordance with a second protocol;

15 from a first server partition comprising a plurality of servers, each of which provides the requested first media processing resource, selecting a first server that is least loaded; and

20 forwarding the first morphed request to the first server in accordance with the second protocol.

2. The method of claim 1, further comprising:

25 receiving a second request from the client requesting the first media processing resource, the second request being formatted in accordance with the first protocol;

morphing the second request into a second morphed request formatted in accordance with the second protocol;

30 if the first server is still least loaded, forwarding the second morphed request to the first server in accordance with the second protocol; and

35 if the first server is not still least loaded, from the first server partition selecting a second server that is least loaded, and forwarding the second morphed request to the second server in accordance with the second protocol.

3. The method of claim 1 or 2, further comprising:

receiving a second request from the client requesting a second media processing resource, the second request being formatted in accordance with the first protocol;

5 morphing the second request into a second morphed request formatted in accordance with the second protocol;

10 from a second server partition comprising a plurality of servers, each of which provides the requested second media processing resource, selecting a second server that is least loaded; and

forwarding the second morphed request to the second server in accordance with the second protocol.

15 4. The method of claim 1, 2 or 3, wherein said step of selecting the first server that is least loaded comprises:

20 applying resource selection rules to at least one pattern contained in a message component to identify the requested media processing resource, the message component selected from the group consisting of a uniform resource identifier (URI) and a content-type header; and

25 referencing at least one partition identifier to identify the server partition comprising the servers providing the requested media processing resource.

5. The method of any of claims 1 to 4, wherein said step of selecting the first server that is least loaded comprises:

30 referencing a routing list containing a list of available servers; and
querying each of the available servers.

6. The method of claim 5, further comprising:

35 removing unavailable servers from the routing list.

7. The method of any of claims 1 to 6, wherein said step of forwarding the morphed request comprises selecting a media servlet on the first server that is associated with the media processing resource.

5 8. The method of claim 7, further comprising establishing a client session with the media servlet or further comprising supporting the client session by passively allowing cookies to bypass load balancing routing.

9. The method of any of claims 1 to 8, further comprising:

10 defining the first protocol to be a Media Resource Control Protocol (MRCP);

15 defining the second protocol to be a Hypertext Transfer Protocol (HTTP); and

distributing the first media processing resource across the plurality of servers.

20 10. A machine readable storage having stored thereon a computer program having a plurality of code sections comprising:

25 code for receiving a first request from a client requesting a first media processing resource, the request being formatted in accordance with a first protocol;

code for morphing the first request into a first morphed request formatted in accordance with a second protocol;

30 code for selecting a first server that is least loaded from a first server partition comprising a plurality of servers, each of which provides the requested first media processing resource; and

35 code for forwarding the first morphed request to the first server in accordance with the second protocol.

11. A system for load balancing media resources among a plurality of servers comprising:

a protocol morpher that receives a first request from a client requesting a first media processing resource, the request being formatted in accordance with a first protocol, morphs the first request into a first morphed request formatted in accordance with a second protocol; and

5

a load balancer that selects a first server that is least loaded from a first server partition comprising a plurality of servers, each of which provides the requested first media processing resource, and forwards the first morphed request to the first server in accordance with the second protocol.

10




FIG. 1

200

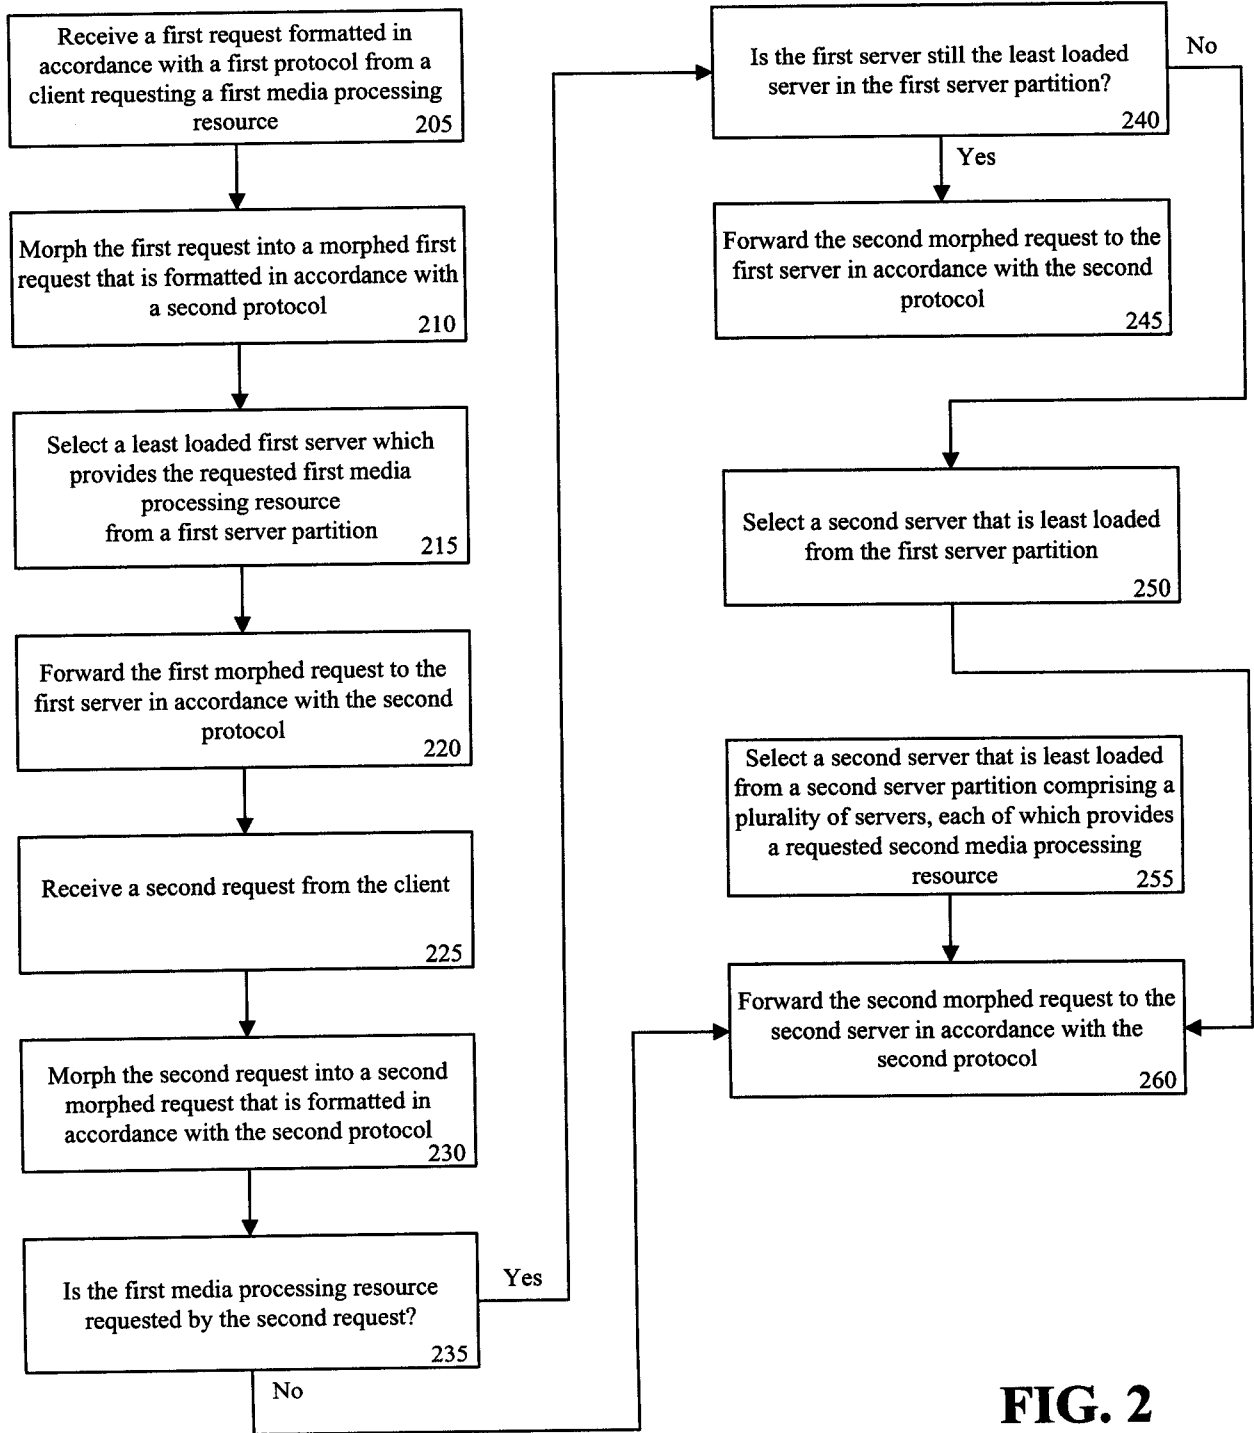


FIG. 2

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2006/069496

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/08

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)
H04L

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, COMPENDEX, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2004/125794 A1 (MARQUETTE BRIAN [US] ET AL) 1 July 2004 (2004-07-01) abstract paragraphs [0002] - [0004], [0007] - [0010], [0014] - [0016], [0028], [0029], [0036] - [0040]	1-11
Y	J. CHAMBERLAIN ET ALL: "IBM WebSphere Voice Server V5.1.1/V5.1.2 and AVAYA Interactive Response V1.3:An Interoperability Guide" INTERNET, November 2005 (2005-11), pages 1-46, XP002424512 Retrieved from the Internet: URL: http://www.redbooks.ibm.com/redpapers/pdfs/redp3977.pdf [retrieved on 2007-03-13] page 10 - page 12	1-11

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier 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.
- *G* document member of the same patent family

Date of the actual completion of the international search

14 March 2007

Date of mailing of the international search report

27/03/2007

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

LOPEZ MONCLUS, I

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2006/069496

<i>C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT</i>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>WO 03/093975 A (BEA SYSTEMS INC [US]) 13 November 2003 (2003-11-13) abstract page 2, line 15 - page 3, line 20 page 10, line 1 - page 11, line 5</p>	1-11
A	<p>SHANMUGHAM CISCO SYSTEMS S ET AL: "Media Resource Control Protocol Version 2 (MRCPv2)" IETF STANDARD-WORKING-DRAFT, INTERNET ENGINEERING TASK FORCE, IETF, CH, vol. speechsc, no. 9, 7 December 2005 (2005-12-07), XP015042729 ISSN: 0000-0004 the whole document</p>	1-11
P,A	<p>J. CHAMBERLAIN ET ALL: "IBM WebSphere Voice Server for Multiplatforms V5.1.1/V5.1.2 Handbook" INTERNET, December 2005 (2005-12), pages 1-94, XP002424511 Retrieved from the Internet: URL:http://www.redbooks.ibm.com/redbooks/pdfs/sg246447.pdf [retrieved on 2007-03-13] page 9 - page 20 page 87 - page 92</p>	1-11

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2006/069496

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004125794 A1	01-07-2004	NONE	
WO 03093975 A	13-11-2003	AU 2003225199 A1 US 2004010611 A1	17-11-2003 15-01-2004