

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 June 2008 (26.06.2008)

PCT

(10) International Publication Number
WO 2008/074668 A1

- (51) International Patent Classification:
G06F 11/34 (2006.01)
- (21) International Application Number:
PCT/EP2007/063538
- (22) International Filing Date:
7 December 2007 (07.12.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/614,220 21 December 2006 (21.12.2006) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, NY 10504 (US).
- (71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; P.O. Box 41, North Harbour, Portsmouth Hampshire PO6 3AU (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **HUNT, Christian**

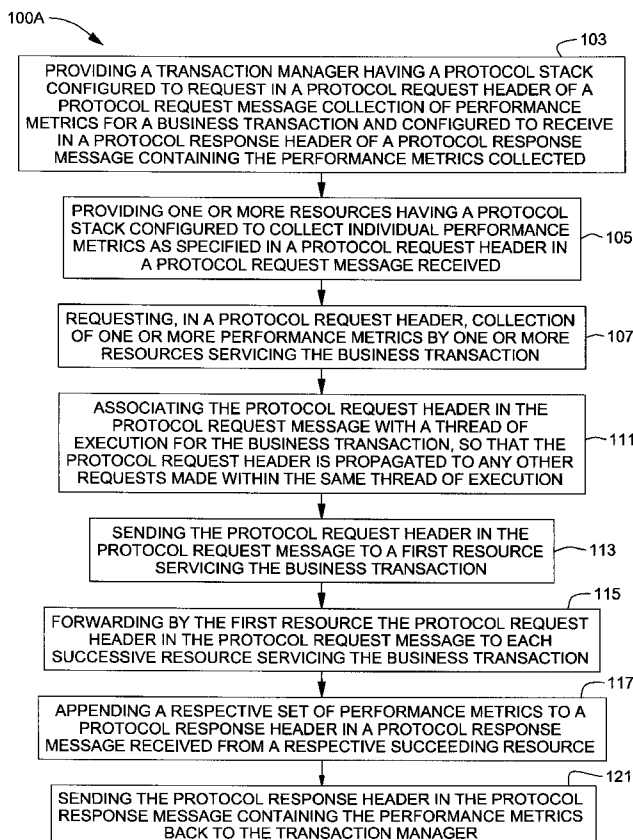
Lee [US/US]; 100 Keithwood Lane, Cary, NC 27511 (US).
MELI, Roger Michael [GB/US]; 1704 West Paces Ferry Road, Raleigh, NC 27613 (US). **NASTACIO, Denilson** [BR/US]; 801 Brickstone Drive, Apex, NC 27502 (US). **SUBRAMANIAN, Balan** [IN/US]; 405 Casey Brook Court, Cary, NC 27519 (US).

(74) Agent: **ROBERTS, Scott**; 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, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, 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.

[Continued on next page]

(54) Title: METHOD, SYSTEM AND PROGRAM PRODUCT FOR MONITORING RESOURCES SERVICING A BUSINESS TRANSACTION



(57) Abstract: A method, system and program product for monitoring resources servicing a business transaction is provided. The method includes providing a transaction manager that monitors a business transaction serviced by one or more resources, with each of the resources and the transaction manager having a protocol stack configured to send and receive load-monitoring headers in protocol messages. The method further includes specifying, in a protocol request header in a protocol request message, collection of performance metrics for a resource, collecting, by the resources, the performance metrics specified in the protocol request message and sending the performance metrics collected to the transaction manager, in a protocol response header in a protocol response message. Further, the method includes preconfiguring with policies the protocol stack on a subset of the resources, such that the preconfigured resources can act as monitoring agents on behalf of peer resources and send alerts to a resource manager.

WO 2008/074668 A1



(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, MT, 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*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

**METHOD, SYSTEM AND PROGRAM PRODUCT FOR MONITORING
RESOURCES SERVICING A BUSINESS TRANSACTION**

5 **FIELD OF THE INVENTION**

The present invention relates to a method, system and computer program product for monitoring resources servicing a business transaction in a service-oriented architecture (SOA). In particular, the present invention relates to a method, system and computer
10 program product for decorating regular SOA exchanges with additional monitoring information, which is carried in the header portion of a response message along with the payload. Further, the invention relates to a method of cross inspecting a header in a response message received by a resource, whereby the resource acts as a monitoring agent on behalf of other peer resources.

15

BACKGROUND OF THE INVENTION

In today's business environment, customers are increasingly conducting online or electronic transactions with organizations and/or businesses, such as retailers, banking services, etc. In
20 a typical scenario, a customer seeking to conduct an online transaction, accesses a business' web page via the Internet and then may have to supply personal information depending on the type of online transaction. For instance, if purchasing an item online, the customer typically has to provide personal information that identifies the customer, such as, name, address, a credit card number, etc. for purchasing the item. Invariably, the data for such
25 online transactions may be handled by multiple computers on a network, where each computer may carry out one aspect of the transaction. Any one of these computers, if overloaded, etc., could lead to delays in processing the transaction, which could negatively impact the customer. As such, there is a need for a business and/or organization to provide a cost-effective way to develop monitoring systems that can track such electronic transactions
30 in order to ensure that customers are kept satisfied.

SUMMARY OF THE INVENTION

In a first aspect of the invention, there is provided a method of monitoring resources servicing a business transaction in a service-oriented architecture. The method includes providing a transaction manager configured to monitor a business transaction serviced by one or more resources in a service-oriented architecture. In an embodiment, each of the one or more resources and the transaction manager has deployed thereon a protocol encoding-decoding component configured to send one or more protocol request messages requesting collection of performance metrics and configured to receive one or more protocol response messages containing the performance metrics collected. The method includes specifying in a protocol request header in a protocol request message collection of one or more performance metrics for at least one resource of the one or more resources servicing the business transaction, collecting the one or more performance metrics specified in the protocol request message and sending to the transaction manager, in a protocol response header in a protocol response message, the one or more performance metrics collected by the at least one resource of the one or more resources servicing the business transaction. In an embodiment, the transaction manager is configured to monitor one or more transaction metrics associated with the business transaction and is configured to analyze the one or more transaction metrics monitored for the business transaction in view of the one or more performance metrics collected from the one or more resources. The method further includes providing a resource manager configured to collect one or more performance metric events from one or more monitoring agents deployed on one or more resources of the one or more resources. The method further includes associating the protocol request header in the protocol request message with a thread of execution for the business transaction, and propagating the protocol request header in the protocol request message to any other requests made within the same thread of execution for the business transaction, wherein the protocol request header in the protocol request message is replicated through all exchanges associated with the business transaction. The method further includes pre-configuring at least a respective protocol encoding-decoding component corresponding to a respective resource of the one or more resources with policies, wherein the respective resource inspects an aggregate value for a performance metric of the one or more performance metrics in a respective protocol response header received from a succeeding resource and wherein the

respective resource sends an alert to the resource manager when the pre-set value for the performance metric is exceeded. In an embodiment, the sending step further includes appending, by a respective preceding resource of the one or more resources, a respective set of performance metrics to a protocol response header in a protocol response message received from a respective succeeding resource of the one or more resources, and sending the protocol response message to the transaction manager. Further, in an embodiment, the sending step includes aggregating in the protocol response header of the protocol response message each of the one or more performance metrics corresponding to a resource of the one or more resources servicing any other requests made within the same thread of execution for the business transaction, and sending the protocol response message to the transaction manager. In an embodiment, the protocol includes Simple Object Access Protocol (SOAP).

In another aspect of the invention, there is provided a system for monitoring resources servicing a business transaction in a service-oriented architecture. The system includes one or more resources configured to service a business transaction, with a plurality of the one or more resources having deployed thereon a protocol stack. The system further includes a transaction manager having deployed thereon the protocol stack and being configured to issue in a protocol request header in a protocol request message to the plurality of the one or more resources to monitor one or more performance metrics for the business transaction serviced by the plurality of the one or more resources, each of the plurality of the one or more resources being configured to collect the one or more performance metrics for the business transaction and to aggregate the one or more performance metrics in a protocol response header in a protocol response message that is sent to the transaction manager, wherein the protocol request header in the protocol request message is associated with a thread of execution for the business transaction and wherein the protocol request header in the protocol request message is propagated to any other requests made within the same thread of execution for the business transaction, such that the protocol request header in the protocol request message is replicated through all exchanges associated with the business transaction. The system further includes a resource manager configured to collect one or more performance metric events from one or more monitoring agents deployed on the one or more resources. In an embodiment, the transaction manager is further configured to monitor one or more transaction metrics for the business transaction serviced by the one or more

resources and is configured to analyze the one or more transaction metrics monitored for the one or more resources servicing the business transaction in view of the one or more performance metrics collected by the one or more resources. In an embodiment, a respective preceding resource of the plurality of the one or more resources is configured to append a
5 respective set of performance metrics to the protocol response header in the protocol response message received from a respective succeeding resource of the one or more resources. In an embodiment, the protocol stack corresponding to one or more of the plurality of the one or more resources is pre-configured with policies in order to compare an aggregate value of a performance metric of the one or more performance metrics received in
10 the protocol response header to a pre-set value for the performance metric, such that the one or more of the plurality of the one or more resources sends an alert to the resource manager when the pre-set value for the performance metric is exceeded. In an embodiment, the protocol stack is a Simple Object Access Protocol (SOAP) stack, the protocol request header is a Simple Object Access Protocol (SOAP) request header, the protocol response header is a
15 Simple Object Access Protocol (SOAP) response header; the protocol request message is a Simple Object Access Protocol (SOAP) request message; and the protocol response message is a Simple Object Access Protocol (SOAP) response message.

In yet another aspect of the invention, there is provided a computer program product for
20 monitoring metrics associated with a business transaction. The computer program product includes a computer readable medium, first program instructions to provide a transaction manager configured to specify, in a request header in a request message issued to one or more resources servicing a business transaction, collection of one or more performance metrics, each of the one or more resources and the transaction manager having deployed
25 thereon a protocol stack. The computer program product further includes second program instructions to collect the one or more performance metrics specified in the request message and third program instructions to send to the transaction manager in a response header in a response message the one or more performance metrics collected by the one or more resources servicing the business transaction. In an embodiment, the second program
30 instructions further include instructions to associate the request header in the request message with a thread of execution for the business transaction and to propagate the request header in the request message to any other request messages made within the thread of

execution for the business transaction. Further, in an embodiment, the second program instructions further include instructions to append a respective set of performance metrics to the response header in the response message received from a succeeding resource of the one or more resources and to send the response message to the transaction manager. The computer program product further includes fourth program instructions to pre-configure the protocol stack deployed on the one or more resources with policies in order to compare an aggregate value of a performance metric of the one or more performance metrics in a respective response header received from a succeeding resource to a pre-set value for the performance metric and to send an alert to a resource manager when the pre-set value for the performance metric is exceeded. In an embodiment, the transaction manager is further configured to monitor one or more transaction metrics associated with the business transaction and is further configured to analyze the one or more transaction metrics monitored for the business transaction in view of the one or more performance metrics collected from the one or more resources. In an embodiment, the protocol stack is a Simple Object Access Protocol (SOAP) stack, the request header is a Simple Object Access Protocol (SOAP) request header, the response header is a Simple Object Access Protocol (SOAP) response header, the request message is a Simple Object Access Protocol (SOAP) request message, and the response message is a Simple Object Access Protocol (SOAP) response message. In an embodiment, each of the first, second, third and fourth program instructions are stored on the computer readable medium.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of the invention:

FIG. 1A depicts a flowchart outlining the steps involved in an embodiment of monitoring resources servicing a business transaction, in accordance with an embodiment of the present invention.

FIG. 2A depicts a flowchart outlining the steps involved in another embodiment of monitoring resources servicing a business transaction, wherein one or more resources

additionally act as monitoring agents on behalf of peer resources, in accordance with an embodiment of the present invention.

FIG. 1B is a schematic block system diagram illustrating an embodiment of a method and system for monitoring resources servicing a business transaction, corresponding to the method outlined in FIG. 1A, in accordance with an embodiment of the present invention.

FIG. 2B is a schematic block system diagram illustrating another embodiment of a method and system for monitoring resources servicing a business transaction, corresponding to the method outline in FIG. 2A, in accordance with an embodiment of the present invention.

FIG. 3 is a schematic block system diagram illustrating yet another embodiment of a method and system for monitoring resources servicing a business transaction in a service-oriented architecture, in accordance with an embodiment of the present invention.

FIG. 4 is a schematic block system diagram illustrating an embodiment of a computer infrastructure for monitoring resources servicing a business transaction in a service-oriented architecture, in accordance with an embodiment of the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

Reference throughout this specification to “one embodiment,” “an embodiment,” or similar language means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases “in one embodiment,” “in an embodiment,” and similar language throughout this specification may, but do not necessarily, all refer to the same embodiment.

Moreover, the described features, structures, or characteristics of the invention may be combined in any suitable manner in one or more embodiments. It will be apparent to those skilled in the art that various modifications and variations can be made to embodiments of the present. Reference will now be made in detail to the preferred embodiments of the invention.

Reference is now made to FIG. 1A and 2A, which depict flowcharts outlining different embodiments of a method of monitoring resource or performance metrics or information

across business method invocation. The term “business method invocation” refers to a simple request from one application to another application, such as “verify credit card number XYZ for a payment of \$70”. As shown in FIG. 1A, in an embodiment 100A, the method of monitoring resource metrics starts with providing in step 103 a transaction manager (discussed further herein below with respect to FIGS. 1B and 2B) that has a protocol stack that is configured to request in a protocol request header of a protocol request message collection of performance metrics for a business transaction and is further configured to receive in a protocol response header of a protocol response message containing the performance metrics collected by the resources. Further, the method includes providing in step 105 one or more resources having a protocol stack configured to collect individual performance metrics as specified in a protocol request header in a protocol request message received by the resources. In step 107, the method includes requesting, in a protocol request header, collection of one or more performance metrics by one or more resources servicing the business transaction and associating, in step 111, the protocol request header in the protocol request message with a thread of execution for the business transaction, so that the protocol request header is propagated to any other requests made within the same thread of execution. Further, the method includes sending, in step 113, the protocol request header in the protocol request message to a first resource servicing the business transaction, forwarding by the first resource, in step 115, the protocol request header in the protocol request message to each successive resource servicing the business transaction. The method further includes appending, in step 117, a respective set of performance metrics to a protocol response header in a protocol response message received from a respective succeeding resource and sending, in step 121, the protocol response header in the protocol response message containing the performance metrics back to the transaction manager.

Further, as shown in FIG. 2A, in another embodiment 200A, the method of monitoring resources servicing a business transaction includes providing a transaction manager, in step 203, that has a protocol stack configured to request in a protocol request header of a protocol request message collection of performance metrics for a business transaction and configured to receive in a protocol response header of a protocol response message containing the performance metrics collected. Further, the method includes providing one or more

resources in step 205 that have a protocol stack configured to collect individual performance metrics as specified in a protocol request header in a protocol request message received, where one or more of the resources have a protocol stack pre-configured with policies or rules and in step 207 requesting, in a protocol request header, collection of one or more performance metrics by one or more resources servicing the business transaction. Further, the method includes in step 209 requesting, in the protocol request header, reporting of any performance metrics collected on behalf of the other resources by the one or more resources having a protocol stack with pre-configured policies when a certain preset value for a performance metric has been exceeded and associating in step 211 the protocol request header in the protocol request message with a thread of execution for the business transaction, so that the protocol request header is propagated to any other requests made within the same thread of execution. Furthermore, the method includes sending, in step 213, the protocol request header in the protocol request message to a first resource servicing the business transaction. In step 215, the method includes forwarding by the first resource the protocol request header in the protocol request message to each successive resource servicing the business transaction. As such, the method includes appending in step 217 a respective set of performance metrics to a protocol response header in a protocol response message received from a respective succeeding resource and inspecting by a resource having a protocol stack pre-configured with policies, in step 219, an aggregate value for a performance metric collected in a respective protocol response header received from a succeeding resource, and sending in step 221, the protocol response header in the protocol response message containing the performance metrics back to the transaction manager. Further, the method includes sending in step 223 an alert to a resource manager if the aggregate value for a performance metric inspected exceeds a pre-set value for the performance metric.

Further, reference is now made to FIGS. 1B and 2B, which depict the various embodiments of the method or process steps carried out by various components involved in a business transaction, as outlined above in FIGS. 1A and 2A herein above. In particular the method outlined above for FIG. 1A corresponds to the embodiment shown in FIG. 1B, whereas, the method outlined above for FIG. 2A corresponds to the embodiment shown in FIG. 2B. Since there are aspects of both embodiments 100B and 200B that are similar, the similar

aspects will be described together with respect to FIGS. 1B and 2B, whereas, separate features will be discussed separately herein below. As shown in FIGS. 1B and 2B, the method of monitoring resources servicing a business transaction comprises providing a transaction manager 106 (206 in FIG. 2B) configured to monitor a business transaction serviced by one or more resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) in a service-oriented architecture (SOA). Business transactions implemented using a Service-Oriented Architecture (SOA) refers to the calls made from a resource providing one web service to another resource providing another web service. A resource may be any device, machine or system that services a business transaction, such as, a web server, a credit card verification system, a database server, an e-mail server, etc. Further, a transaction manager 106 (206 in FIG. 2B) is an application or system that monitors or collects transaction metrics or transaction data associated with the business transaction with respect to what an end user is experiencing. In particular, the transaction manager 106 (206 in FIG. 2B) may track transaction metrics, such as, how much time a business transaction took from an end user's perspective and/or how much time was spent by each resource 110, 114, 118 and/or 122 (210, 214, 218 and/or 222 in FIG. 2B) that processed a business transaction, etc. The transaction manager 106 (206 in FIG. 2B) tracks such transaction metrics or transaction information and may send a breakdown of time spent by each resource to an administrator at a remote location, so that the administrator can take appropriate action, as necessary. However, as disclosed in the method embodiments of FIGS. 1B and 2B, the transaction manager 106 (206 in FIG. 2B) is configured to track performance or resource or resource-level metrics for the one or more resources servicing a business transaction, such that the transaction manager 106 (206 in FIG. 2B) can interpret the transaction metrics that are tracked in view of the performance metrics for the individual resources that are gathered, as discussed herein below. FIGS. 1B and 2B show a business transaction that is serviced (indicated with solid arrows 112, 116, 120, 124, 126 and 128 in FIG. 1B; 212, 216, 220, 224, 226 and 228 in FIG. 2B) by four resources, namely, resource 1 (reference numerals 110 in FIG. 1B and 210 in FIG. 2B), resource 2 (reference numerals 114 in FIG. 1B and 214 in FIG. 2B), resource 3 (reference numerals 118 in FIG. 1B and 218 in FIG. 2B) and resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B). For instance, resource 1 (reference numerals 110 in FIG. 1B and 210 in FIG. 2B) may be a web server and the web server 110 (210 in FIG. 2B) calls (arrow 112 in FIG. 1B and arrow 212 in FIG. 2B) resource 2

(reference numeral 114 in FIG. 1B, 214 in FIG. 2B) for a web service and resource 2 calls (arrow 116 in FIG. 1B and 216 in FIG. 2B) resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) for another web service and resource 3 calls (arrow 120 in FIG. 1B and 220 in FIG. 2B) resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) for yet another web service. For instance, resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B) may be a “shopping cart checkout” server that a customer accesses when placing an item in a shopping cart, whereas, resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) may be a “product in-stock verification” server that checks to see whether or not the item is in stock, whereas, resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) may be a “credit card verification” server that verifies the customer’s credit card information supplied for the business transaction. Alternatively, for instance, resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B) may call a shipment department (not shown in FIGS. 1B and 2B) to check whether or not an item or product is in inventory and then resource 2 calls (arrow 120 in FIG. 1B and 220 in FIG. 2B) resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) and then resource 3 calls (arrow 120 in FIG. 1B and 220 in FIG. 2B) resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B).

Although, the business transaction flow is shown in a linear manner in FIGS. 1B and 2B, as shown in FIG. 3, the business transaction flow may not be linear in that there may be other steps or processes that are carried out simultaneously by a resource. As shown in the embodiment 300 in FIG. 3, the transaction manager 306 sends a request message to resource 1 (reference numeral 310), which in turn, sends messages to resource 2 (reference numeral 314), resource 3 (reference numeral 318) and resource 4 (reference numeral 322) simultaneously. For instance, resource 1 (reference numeral 310) may be a web server, that contacts resource 2 (reference numeral 314), which may be a database server for checking inventory, whereas, resource 3 (reference numeral 318) may be a credit card verification server and resource 4 (reference numeral 322) may be an e-mail server that sends a message regarding the business transaction. Accordingly, each of the resources, resource 4 (reference numeral 322), resource 3 (reference numeral 318) and resource 2 (reference numeral 314) sends its own response back to resource 1 (reference numeral 310), which sends the response back to the transaction manager 306. FIG. 3 is used here only to show an example of a non-linear transaction and will not be discussed in detail given that the teachings of FIG. 1B and

2 are applicable to FIG. 3 as well. Going back to FIGS. 1B and 2B, in an embodiment, each of the four resources, 110 (210 in FIG. 2B), 114 (214 in FIG. 2B), 118 (218 in FIG. 2B), and/or 122 (222 in FIG. 2B), and the transaction manager 106 (206 in FIG. 2B) has deployed thereon an instrumented protocol encoding-decoding component or an instrumented protocol stack. Alternatively, the transaction manager 106 (206 in FIG. 2B) and only a subset of the four resources, 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) may have an instrumented protocol. In an embodiment, the transport protocol for business transactions implemented using a service-oriented architecture (SOA) is SOAP (Simple Object Access Protocol). The SOAP specification is fully described at <http://www.w3.org/TR/soap/>. A “protocol stack” or “SOAP stack” or an “encoding-decoding component” refers to the layer of software used to encode and decode protocol messages, in particular, SOAP messages. A SOAP message is a text message using a markup language (XML –Extensible Markup Language), where the message has distinct parts: the body or payload and the header. A SOAP message contains one body but may contain one or more headers. A header in a message contains information intended for the transmission nodes between the caller and the respondent, including themselves, such as “return the response to this address” or “this message is part of a larger transaction number 987”. As such, each header in a SOAP message may serve a completely different purpose, such as defining security permissions, transaction markers, delivery parameters, etc. The body or payload contains the message intended from the caller to the respondent, such as “verify credit card number XYZ for a payment of \$70” or the response from the respondent to the caller, such as “credit approved”. In particular, an “instrumented protocol stack” refers to a protocol stack that has the addition functionality, which allows the protocol stack (the layer of software that knows how to encode and decode SOAP messages) to understand the headers (also referred to as “load-monitoring” headers) that request resource-level data or information or performance metrics, in addition to a request carried in the payload of the message, as described herein above with respect to FIGS. 1A and 2A. As such, any other exchange protocol that implements the same concepts of body/header separation in message exchanges, such as SOAP, can be used for practicing the methods of preferred embodiments of the present invention.

Referring to FIGS. 1B and 2B, the transaction manager 106 (206 in FIG. 2B) having an instrumented SOAP stack, issues a SOAP request message 108 (208 in FIG. 2B) to resource 1, reference numeral 110 (210 in FIG. 2B). In particular, the SOAP request message 108 (208 in FIG. 2B) contains a SOAP header that requests collection of one or more performance or resource metrics by one or more of the four resources 110, 114, 118 and 122 (210, 214, 218 and 222) servicing the business transaction, in addition, to any other requests that may be carried in other headers and the actual request carried in the body or payload of the SOAP request message 108 (208 in FIG. 2B). Some examples of performance metrics that can be collected for a resource may include Central Processing Unit (CPU) usage, memory paging, disk usage, disk freespace, disk fragmentation, memory utilization, thread counts, number of processes, bytes read (network I/O), bytes written (network I/O), etc. In particular, as shown in FIGS. 1B AND 2B, the request information for collecting the one or more performance metrics is specified and carried out in a protocol request header (SOAP request header) of a protocol request message (SOAP request message) is forwarded (solid arrows 108, 112, 116 and 120; solid arrows 208, 212, 216 and 220 in FIG. 2B) to each of the resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B), whereas, the performance metrics collected by the one or more resources servicing the business transaction is sent back (solid arrows 124, 126, 128 and 130; solid arrows 224, 226, 228 and 230) to the transaction manager 106 (206 in FIG. 2B) in a protocol response header (SOAP response header) in a protocol response message (SOAP response message). Moreover, each of the resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) associates the protocol (SOAP) request header in the protocol (SOAP) request message with a thread of execution for the business transaction, thus, propagating the protocol (SOAP) request header in the protocol (SOAP) request message to any other requests made within the same thread of execution for the business transaction.

As mentioned herein above, preferably each of the four resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) as well as the transaction manager 106 (206 in FIG. 2B) has an instrumented SOAP stack, so that as the SOAP request message is issued (arrow 108 in FIG. 1B; arrow 208 in FIG. 2B) from the transaction manager 106 (206 in FIG. 2B) to resource 1 (reference numeral 110 in FIG. 1B and 210 in FIG. 2B), the instrumented SOAP stack on resource 1 is able to follow the instructions in the request header in the SOAP

request message 108 (208 in FIG. 2B). Accordingly, resource 1 collects the performance metrics that are requested by the transaction manager 106 (206 in FIG. 2B) in the SOAP header of the SOAP request message 108 (208 in FIG. 2B). Similarly, as the request message 108 (208 in FIG. 2B) is forwarded, resource 1 (reference numeral 110) propagates the request header contained in the SOAP request message 108 (208 in FIG. 2B) and sends it (arrow 112 in FIG. 1B; arrow 212 in FIG. 2B) to resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B), such that, the instrumented SOAP stack on resource 2 collects its own performance metrics for resource 2 that are requested by the transaction manager 106 (206 in FIG. 2B) in the SOAP header of the SOAP request message 108 (208 in FIG. 2B). Again, as the message 108 (208 in FIG. 2B) is forwarded from resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B) to resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B), resource 2 propagates the request header contained in the SOAP request message 108 (208 in FIG. 2B) and sends it (arrow 116 in FIG. 1B; arrow 216 in FIG. 2B) to resource 3, so that the instrumented SOAP stack on resource 3 collects its own performance metrics for resource 3 that are requested by the transaction manager 106 (206 in FIG. 2B) in the SOAP header of the SOAP request message 108 (208 in FIG. 2B). Moreover, as the message 108 (208 in FIG. 2B) is forwarded from resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) to resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B), resource 3 propagates the request header contained in the SOAP request message 108 (208 in FIG. 2B) and sends it (arrow 120 in FIG. 1B; arrow 220 in FIG. 2B) to resource 4, such that the instrumented SOAP stack on resource 4 collects its own performance metrics that are requested by the transaction manager 106 (206 in FIG. 2B) in the SOAP header of the SOAP request message 108 (208 in FIG. 2B). As a result, the protocol (SOAP) request header in the protocol (SOAP) request message is replicated through all exchanges or messages associated with the business transaction.

In an embodiment, the instrumented SOAP stacks on each of the resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) collects performance metric or resource metric information and adds or appends the information collected to SOAP response headers in a response message pertaining to the business transaction request. In particular, the performance metrics collected by resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) is added to a SOAP response header in a SOAP response message that is sent back

(arrow 124 in FIG. 1B; arrow 224 in FIG. 2B) from resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) to resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B). In turn, resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) adds the performance metrics collected by resource 3 to the SOAP response header in the SOAP response message received from resource 4 and sends it back (arrow 126 in FIG. 1B; arrow 226 in FIG. 2B) to resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B), which in turn adds the performance metrics collected by it to the SOAP response header in the SOAP response message received from resource 3 and sends it back (arrow 128) to resource 1 (reference numeral 110 in FIG. 1B and 210 in FIG. 2B). Again, resource 1 adds performance metrics collected by it to the SOAP response header in the SOAP response message received from resource 2 and sends (arrows 130 in FIG. 1B; arrow 230 in FIG. 2B) the SOAP response message back to the transaction manager 106 (206 in FIG. 2B). Further, the instrumented SOAP stack aggregates, in the SOAP response header of the SOAP response message sent to the transaction manager 106 (206 in FIG. 2B), each of the one or more performance metrics collected for a resource servicing any other requests made within the same thread of execution for the business transaction before sending the SOAP response message to the transaction manager. Moreover, since the instrumented SOAP stack on each resource is responsible for collecting metrics for the resource, it is preferable that each resource has an instrumented SOAP stack that can collect performance metrics, but it is not necessary. So if each of the resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) has installed thereon an instrumented SOAP stack, the instrumented SOAP stack on a preceding resource can add its own performance metrics to the response header they received from a successive resource they called. Accordingly, the final response header in the SOAP response message sent (solid arrow 130 in FIG. 1B; solid arrow 230 in FIG. 2B) to the transaction manager 106 (206 in FIG. 2B) contains one response header with the performance metrics for each of the resources 110, 114, 118 and 122 (210, 214, 218 and 222 in FIG. 2B) contributing to a comprehensive answer. However, if a resource has a SOAP stack that is non-instrumented (a non-instrumented SOAP stack), the non-instrumented SOAP stack simply passes the performance or resource load-monitoring headers to a successive resource without collecting its own performance metrics. Hence, the performance metrics for a non-instrumented SOAP stack on a resource will not be included in the final response header in the response message delivered to the transaction manager.

In an embodiment, the method further comprises providing a resource manager 102 (202 in FIG. 2B) that is configured to collect one or more performance metric events from one or more monitoring agents deployed on one or more of the resources 110, 114, 118 and/or 122 (210, 214, 218 and/or 222 in FIG. 2B) in order to manage the resources. A resource manager 102 (202 in FIG. 2B) is an application that is configured to receive events from a monitoring agent running on a resource. Further, a monitoring agent is a program that is run on one or more resources, whose task is to monitor performance or resource metrics on those resources, such as, CPU data, disk usage, etc. If the monitoring agent observes something that is problematic, for example, that the CPU on a resource has reached a 70% threshold, the monitoring agent generates an event (also referred to as performance metric event) and sends it to the resource manager 102 (202 in FIG. 2B). For instance, as shown in FIGS. 1B AND 2B, resource 1 (reference numeral 110 in FIG. 1B and 210 in FIG. 2B) has a monitoring agent 134 (234 in FIG. 2B) that collects performance metrics on resource 1. The performance metrics collected by the monitoring agent 134 (234 in FIG. 2B) may be different than the one requested by the transaction manager 106 (206 in FIG. 2B). The monitoring agent 134 (234 in FIG. 2B) sends (dotted arrow 138 in FIG. 1B; dotted arrow 238 in FIG. 2B) the performance metrics collected for resource 1 to a resource manager 102 (202 in FIG. 2B). The resource manager 102 (202 in FIG. 2B) is configured to receive alerts (dotted arrow 138 in FIG. 1B; arrow 238 in FIG. 2B) from the monitoring agent 134 (234 in FIG. 2B) running on resource 1 (reference numeral 110 in FIG. 1B and 210 in FIG. 2B) when a performance metric for resource 1 has exceeded a pre-set value, so that the resource manager 102 (202 in FIG. 2B) may take any appropriate actions needed. Similarly, resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B) has a monitoring agent 136 (236 in FIG. 2B) that collects performance metrics for resource 2. Again, the monitoring agent 136 (236 in FIG. 2B) sends (dotted arrow 140 in FIG. 1B; dotted arrow 240 in FIG. 2B) to the resource manager 102 (202 in FIG. 2B) the performance metrics collected and/or alerts when a performance metric for resource 2 has exceeded a pre-set value or threshold. As such, the resource manager 102 (202 in FIG. 2B) collects all the resource or performance metrics events for the individual resources 110 and 114 (resources 210 and 214 in FIG. 2B) that have a monitoring agent 134 and 136 (234 and 236 in FIG. 2B) running on them, respectively. However, since resource 3 (reference numeral 118 in FIG. 1B and 218 in FIG. 2B) and resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) do not have any

monitoring agents on them, only the information collected by the monitoring agents 134 and 136 (234 and 236 in FIG. 2B) running on resource 1 (reference numeral 110 in FIG. 1B and 210 in FIG. 2B) and resource 2 (reference numeral 114 in FIG. 1B and 214 in FIG. 2B), respectively, gets passed or communicated to the resource manager 102 (202 in FIG. 2B).
5 This means that if, for example, resource 4 (reference numeral 122 in FIG. 1B and 222 in FIG. 2B) is slow (shown by the symbol 132 in FIG. 1B and 232 in FIG. 2B) due to a performance metric, for instance, CPU usage, that has been exceeded, this performance information may not get communicated to the resource manager 102 (202 in FIG. 2B), thus, not addressing the problem with respect to resource 4. Further, the resource manager 102
10 (202 in FIG. 2B) does not actively communicate the information to the transaction manager 106 (206 in FIG. 2B), but the information from the resource manager 102 (202 in FIG. 2B) may be communicated (double-sided arrow 104 in FIG. 1B, double-sided arrow 204 in FIG. 2B) to the transaction manager 106 (206 in FIG. 2B), if requested by the transaction manager 106 (206 in FIG. 2B).

15

In another embodiment, as outlined in FIG. 2A and shown in FIG. 2B, provided is a method 200A and 200B of monitoring a business transaction serviced by one or more resources. , which includes pre-configuring with policies or rules at least a respective instrumented protocol encoding-decoding component (instrumented SOAP stack) corresponding to a
20 respective resource of the one or more resources, such that the respective resource inspects an aggregate value for a performance metric in a respective protocol (SOAP) response header received from a succeeding resource. Further, the respective resource sends an alert to the resource manager when the pre-set value for the performance metric for the succeeding resource is exceeded. Accordingly, the instrumented SOAP stack that is pre-
25 configured with policies or rules further acts as a monitoring agent and is able to send an alert to the resource manager on behalf of a peer resource when a determination is made that a pre-set value or pre-configured threshold for a performance metric on the peer resource has been exceeded. An alert may be implemented as a simple message that is sent directly to the resource manager indicating which metric has exceeded a threshold or pre-set value,
30 probably containing the measured value and the pre-configured threshold or pre-set value. For instance, as shown in FIG. 2B, resource 2 (reference numeral 214) has an instrumented SOAP stack that is also pre-configured with policies (symbol 242 in FIG. 2B) to act as a

monitoring agent on behalf of resource 3 (reference numeral 218) and/or resource 4 (reference numeral 222) in FIG. 2B. In particular, if certain criteria or values have been pre-set for either resource 3 and resource 4 or both resources 3 and 4 in FIG. 2B, then the instrumented and pre-configured SOAP stack on resource 2 (reference numeral 214) can

5 compare the performance metrics collected and received from resource 4 (reference numeral 222) and/or resource 3 (reference numeral 218) and can determine whether or not a particular pre-set threshold value or criterion has been met for either resource 3 and/or resource 4 in FIG. 2B, so that resource 2 (reference numeral 214) can send out an alert (dotted arrow 244) to the resource manager 202 on behalf of the resource (be it resource 3

10 (reference numeral 218) and/or resource 4 (reference numeral 222)), thus, alerting the resource manager 202 in FIG. 2B of the met pre-set threshold values and/or criteria, unlike the situation described earlier for FIG. 1B. Accordingly, in an embodiment, only a few of the instrumented SOAP stacks are pre-configured with policies to act as monitoring agents to report events on behalf of other resources, thus, minimizing the number of monitoring agents

15 required to monitor the resources throughout the network and minimizing any connectivity issues of allowing each and every single resource to communicate their status to a central monitoring infrastructure in addition to communicating to peers regarding regular business transactions. Further, in an embodiment, the resource monitoring infrastructure or resource manager 202 provides a graphical user interface, where the criteria or values for one or more

20 performance metrics can be pre-set. Accordingly, when a transaction manager 206 requests monitoring information, it may ask for both specific performance metrics and for performance metrics when a certain pre-set threshold value has been exceeded. Further, a monitoring agent, such as monitoring agent 236, may still be installed on a resource (for instance, resource 2, reference numeral 214 in FIG. 2B) even if the instrumented SOAP

25 stack on the resource 2, reference numeral 214, is pre-configured with policies (symbol 242) in order to compare and send alerts to the resource manager 202. For instance, a monitoring agent 236 may monitor a different set of performance or resource metrics than those gathered by the pre-configured and instrumented SOAP stack 242 on resources 2, reference numeral 214.

30

In another embodiment, as shown in FIGS. 1B and 2B, provided are different embodiments of a system 100B and 200B for monitoring resources servicing a business transaction in a

service-oriented architecture, as described herein above. The system comprises one or more resources (110, 114, 118 and 122 in FIGS. 1B and 210, 214, 218 and 222 in FIGS. 2B) configured to service a business transaction, with one or more of the resources have deployed thereon an instrumented protocol stack, preferably, each of the resources having
5 deployed thereon an instrumented protocol stack. The system further comprises a transaction manager (106 in FIGS. 1B and 206 in FIGS. 2B) having deployed thereon the instrumented protocol stack and being configured to issue in a protocol request header in a protocol request message sent (solid arrow 108 in FIGS. 1B and 208 in FIGS. 2B) to the four resources in FIGS. 1B and 2B to monitor one or more performance metrics for the business
10 transaction serviced by the four resources. Further, in an embodiment, each of the four resources (shown in FIGS. 1B and 2B) is configured to collect the one or more performance metrics for the business transaction and to aggregate the one or more performance metrics in a protocol response header in a protocol response message that is sent to the transaction manager, where the protocol request header in the protocol request message is associated
15 with a thread of execution for the business transaction, such that the protocol request header in the protocol request message is propagated to any other requests made within the same thread of execution for the business transaction. Accordingly, the protocol request header in the protocol request message is replicated through all exchanges associated with the business transaction, as described in the method hereinabove. The system further comprises a
20 resource manager (102 in FIGS. 1B and 202 in FIGS. 2B) configured to collect one or more performance metric events from one or more monitoring agents (134 and 136 in FIGS. 1B and 234 and 236 in FIGS. 2B) deployed on the one or more resources, in particular, resources 1 and 2 as shown in FIGS. 1B and 2B. In an embodiment, the transaction manager (106 in FIGS. 1B and 206 in FIGS. 2B) is further configured to monitor one or more
25 transaction metrics for the business transaction serviced by the one or more resources (110, 114, 118 and 122 in FIGS. 1B and 210, 214, 218 and 222 in FIGS. 2B) and is configured to analyze the one or more transaction metrics monitored for the one or more resources servicing the business transaction in view of the one or more performance metrics collected by the one or more resources, as described in the method hereinabove. In an embodiment, a
30 respective preceding resource of the plurality of the one or more resources, for instance, resource 3 (118 in FIGS. 1B and 218 in FIGS. 2B) is configured to append a respective set of performance metrics to the protocol response header in the protocol response message

received from a respective succeeding resource of the one or more resources, for instance, resource 4 (122 in FIGS. 1B and 222 in FIGS. 2B). In an embodiment, as shown in FIGS. 2B, the instrumented protocol stack corresponding to a subset of the one or more resources, for instance, only resource 2 (214 in FIGS. 2B) is pre-configured with policies (symbol 242) in order to compare an aggregate value of a performance metric in the protocol response header of a protocol response message received to a pre-set value for the performance metric, such that the pre-configured resource (for instance, resource 2, reference numeral 214) or resources send an alert to the resource manager 202 when the pre-set value for the performance metric is exceeded. In an embodiment, the instrumented protocol stack comprises an instrumented Simple Object Access Protocol (SOAP) stack, wherein the protocol request header comprises a Simple Object Access Protocol (SOAP) request header, wherein the protocol response header comprises a Simple Object Access Protocol (SOAP) response header, wherein the protocol request message comprises a Simple Object Access Protocol (SOAP) request message and wherein the protocol response message comprises a Simple Object Access Protocol (SOAP) response message.

In yet another embodiment, a computer program product is provided for monitoring metrics associated with a business transaction. Preferably, the computer program product is in a form accessible from the computer-usable or computer-readable medium, which provides program codes or instructions for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer-usable or computer readable medium can be any apparatus that can contain, store, communicate, propagate, or transport the codes or instructions for use by or in connection with the instruction execution system, apparatus, or device. Preferably, the medium can include an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. More preferably, the computer-readable medium can include a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read-only memory (ROM), a rigid magnetic disk and an optical disk. Further, examples of optical disks include compact disc – read only memory (CD-ROM), compact disc – read/write (CD-R/W) and digital versatile/video disc (DVD). Embodiments of the invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software

elements. A preferred embodiment is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

5 The computer program product further comprises first program instructions to provide a transaction manager configured to specify, in a request header in a request message issued to one or more resources servicing a business transaction, collection of one or more performance metrics, each of the one or more resources and the transaction manager having deployed thereon an instrumented protocol stack. The computer program product further comprises second program instructions to collect the one or more performance metrics
10 specified in the request message and third program instructions to send to the transaction manager in a response header in a response message the one or more performance metrics collected by the one or more resources servicing the business transaction. In an embodiment, the second program instructions further include instructions to associate the request header in the request message with a thread of execution for the business transaction
15 and to propagate the request header in the request message to any other request messages made within the thread of execution for the business transaction. Further, in an embodiment, the second program instructions further include instructions to append a respective set of performance metrics to the response header in the response message received from a succeeding resource of the one or more resources and to send the response message to the
20 transaction manager. The computer program product further comprises fourth program instructions to pre-configure the instrumented protocol stack deployed on the one or more resources with policies in order to compare an aggregate value of a performance metric of the one or more performance metrics in a respective response header received from a succeeding resource to a pre-set value for the performance metric and to send an alert to a
25 resource manager when the pre-set value for the performance metric is exceeded. In an embodiment, the transaction manager is further configured to monitor one or more transaction metrics associated with the business transaction and is further configured to analyze the one or more transaction metrics monitored for the business transaction in view of the one or more performance metrics collected from the one or more resources. In an
30 embodiment, the instrumented protocol stack comprises an instrumented Simple Object Access Protocol (SOAP) stack, and the request header comprises a Simple Object Access Protocol (SOAP) request header. Further, the response header comprises a Simple Object

Access Protocol (SOAP) response header, the request message comprises a Simple Object Access Protocol (SOAP) request message and the response message comprises a Simple Object Access Protocol (SOAP) response message. In an embodiment, each of the first, second, third and fourth program instructions are stored on the computer readable medium.

5

Referring now to Fig. 4, there is illustrated a system 400 for determining an optimal configuration and associated operational costs for implementing a capacity management service, according to embodiments of the present invention. As depicted, system 400 includes a computer infrastructure 402, which is intended to represent any type of computer architecture that is maintained in a secure environment (i.e., for which access control is enforced). As shown, infrastructure 402 includes a computer system 404 that typically represents a server or the like. It should be understood, however, that although not shown, other hardware and software components (e.g., additional computer systems, such as, application servers, administrative servers, routers, firewalls, etc.) could be included in infrastructure 402.

10

15

In general, a user (any one of user 1, reference numeral 432 through user N, reference numeral 434) may interface with infrastructure 402 for performing a business transaction.

20

Similarly, any an administrator 436 and/or service provider 444 can interface with infrastructure 402 for accessing the system 404 in order to monitor resources, such as resource 1, reference numeral 440 through resource N, reference numeral 442 servicing a business transaction. Further, resource 1 (reference numeral 440) through resource N (reference numeral 442) may interface with infrastructure 402 for servicing the business transaction. As shown, resource 1 (reference numeral 440) has an instrumented protocol stack 441 and resource N (reference numeral 442) has an instrumented protocol stack 443 deployed thereon. Further, the transaction manager 414 also has an instrumented protocol stack 416 deployed thereon. As such, each of the resources, preferably, has an instrumented protocol stack for communicating with the transaction manager 414. Furthermore, the resource manager 430 can also interface with computer system 404 for providing the transaction manager 414 any information that may be requested regarding any resource metric collected by the resource manager through monitoring agents and/or resources acting as monitoring agents. In general, the parties could access infrastructure 402 directly, or over

25

30

a network via interfaces (e.g., client web browsers) loaded on computerized devices (e.g., personal computers, laptops, handheld devices, etc.). In the case of the latter, the network can be any type of network such as the Internet or can be any other network, such as, a local area network (LAN), a wide area network (WAN), a virtual private network (VPN), etc. In
5 any event, communication with infrastructure 402 could occur via a direct hardwired connection (e.g., serial port), or via an addressable connection that may utilize any combination of wire line and/or wireless transmission methods. Moreover, conventional network connectivity, such as Token Ring, Ethernet, WiFi or other conventional communications standards could be used. Still yet, connectivity could be provided by
10 conventional TCP/IP sockets-based protocol. In this instance, the parties could utilize an Internet service provider to establish connectivity to infrastructure 402. It should be understood that infrastructure 402 could be owned and/or operated by a party such as provider 444, or by an independent entity. Regardless, use of infrastructure 402 and the teachings described herein could be offered to the parties on a subscription or fee-basis. In
15 either scenario, an administrator 436 could support and configure infrastructure 402.

Computer system or server 404 is shown to include a CPU (hereinafter “processing unit 406”), a memory 412, a bus 410, and input/output (I/O) interfaces 408. Further, computer system 400 is shown in communication with external I/O devices 424 and storage system
20 422. In an embodiment as shown, the storage system 422 includes policies or rules 426 that may be used to pre-configure the respective instrumented protocol stacks (441 through 443) on any of the respective resources 1 through N (reference numerals 440 and 442), such that the pre-configured resource may additionally act as a monitoring agent. In general, processing unit 406 executes computer program codes, such as the transaction manager
25 application 414, which includes an instrumented protocol stack 416, in an embodiment, an instrumented SOAP stack. While executing the transaction manager application 414, the processing unit 406 can read and/or write data, to/from memory 412, storage system 422, and/or I/O interfaces 408. Bus 410 provides a communication link between each of the components in computer system 400. External devices 424 can include any devices (e.g.,
30 keyboard, pointing device, display, etc.) that enable a user to interact with computer system 400 and/or any devices (e.g., network card, modem, etc.) that enable computer system 400 to communicate with one or more other computing devices.

Computer infrastructure 402 is only illustrative of various types of computer infrastructures for implementing embodiments of the invention. For example, in one embodiment, computer infrastructure 402 includes two or more computing devices (e.g., a server cluster) that communicate over a network to perform the various process steps. Moreover, computer system 400 is only representative of various possible computer systems that can include numerous combinations of hardware. To this extent, in other embodiments, computer system 400 can include any specific purpose computing article of manufacture comprising hardware and/or computer program code for performing specific functions, any computing article of manufacture that includes a combination of specific purpose and general purpose hardware/software, or the like. In each case, the program code and hardware can be created using standard programming and engineering techniques, respectively. Moreover, processing unit 406 may include a single processing unit, or be distributed across one or more processing units in one or more locations, e.g., on a client and server. Similarly, memory 412 and/or storage system 422 can include any combination of various types of data storage and/or transmission media that reside at one or more physical locations. Further, I/O interfaces 408 can include any system for exchanging information with one or more external devices 424. Still further, it is understood that one or more additional components (e.g., system software, math co-processing unit, etc., not shown in Fig. 4) can be included in computer system 400. Similarly, it is understood that the one or more external devices 424 (e.g., a display) and/or storage system(s) 422 could be contained within computer system 404, and not externally as shown.

Storage system 422 can be any type of system (e.g., a database) capable of storing information or data, such as policies 426 that provide rules, such as, what a pre-configured protocol stack on a resource should do when a pre-set resource metric has been exceeded. To this extent, storage system 422 could include one or more storage devices, such as a magnetic disk drive or an optical disk drive. In another embodiment, storage system 422 includes data distributed across, for example, a local area network (LAN), wide area network (WAN) or a storage area network (SAN) (not shown). Although not shown, additional components, such as cache memory, communication systems, system software, etc., may be incorporated into computer system 400.

Accordingly, in an embodiment, the transaction manager 414 analyzes the transaction metrics monitored for the business transaction from an end user's perspective (for instance, any of user 1, reference numeral 432 through user N, reference numeral 434) in view of the performance metrics collected from the resource 1 (reference numeral 440) through resource
5 N (reference numeral 442) and tries to make sense of the performance information along with any transaction metrics it may have collected for the resources 1 through N. For instance, if the CPU on a resource is maxed out but the resource is still meeting a desired target time for completing a transaction request, the performance metric for that resource may be considered acceptable. Further, the transaction manager 414 may obtain information
10 from the resource manager 430, for instance, when a resource servicing a particular transaction does not meet a desired target time. In particular, the transaction manager 414 may be able to figure out a reason based on what performance metrics or data are conveyed by the resource manager 430 (that is, through data obtained through the monitoring agents or through resources that are pre-configured with policies to act as monitoring agents). For
15 instance, the transaction manager 414 may figure out that the reason why a transaction did not meet a desired target time was due to hard disk failure or excessive disk usage, etc. Furthermore, the data collected by the resource manager 430 and the transaction manager 414 may be passed on to an administrator 436 and/or provider 444, thus, providing the administrator 436 and/or provider 444 with a better more comprehensive understanding of
20 the resources.

The foregoing descriptions of specific embodiments of the present invention have been presented for the purpose of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and obviously many
25 modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated.

CLAIMS

1. A method of monitoring resources servicing a business transaction in a service
5 oriented architecture, said method comprising the steps of:

providing a transaction manager configured to monitor a business transaction
serviced by one or more resources in a service oriented architecture, each of said one or
more resources and said transaction manager having deployed thereon a protocol encoding-
decoding component configured to send one or more protocol request messages requesting
10 collection of performance metrics and configured to receive one or more protocol response
messages containing said performance metrics collected;

specifying, in a protocol request header in a protocol request message of said one or
more protocol request messages, collection of one or more performance metrics for at least
one resource of said one or more resources servicing said business transaction;

15 collecting said one or more performance metrics specified in said protocol request
message; and

sending to said transaction manager, in a protocol response header in a protocol
response message, said one or more performance metrics collected by said at least one
resource of said one or more resources servicing said business transaction.

2. The method according to claim 1, wherein said providing step further comprises the
steps of:

providing said transaction manager configured to monitor one or more transaction
metrics associated with said business transaction; and

25 providing said transaction manager configured to analyze said one or more
transaction metrics monitored for said business transaction in view of said one or more
performance metrics collected from said one or more resources.

3. The method according to claim 2, wherein said providing step further comprises the
30 step of:

providing a resource manager configured to collect one or more performance metric events from one or more monitoring agents deployed on one or more resources of said one or more resources.

5 4. The method according to claim 3, wherein said specifying step further includes the steps of:

 associating said protocol request header in said protocol request message with a thread of execution for said business transaction; and

 propagating said protocol request header in said protocol request message to any
10 other requests made within said same thread of execution for said business transaction; wherein said protocol request header in said protocol request message is replicated through all exchanges associated with said business transaction.

15 5. The method according to claim 4, wherein said specifying step further comprises the step of:

 pre-configuring at least a respective protocol encoding-decoding component corresponding to a respective resource of said one or more resources with policies, wherein said respective resource inspects an aggregate value for a performance metric of said one or more performance metrics in a respective protocol response header received from a
20 succeeding resource and wherein said respective resource sends an alert to said resource manager when said pre-set value for said performance metric is exceeded.

25 6. The method according to claim 4, wherein said sending step further comprises the steps of:

 appending, by a respective preceding resource of said one or more resources, a respective set of performance metrics to a protocol response header in a protocol response message received from a respective succeeding resource of said one or more resources; and
 sending said protocol response message to said transaction manager.

30 7. The method according to claim 5, wherein said sending step further comprises the steps of:

aggregating in said protocol response header of said protocol response message each of said one or more performance metrics corresponding to a resource of said one or more resources servicing said any other requests made within said same thread of execution for said business transaction; and .

5 sending said protocol response message to said transaction manager.

8. The method according to claim 7, wherein said protocol comprises Simple Object Access Protocol (SOAP).

10 9. A system for monitoring resources servicing a business transaction in a service oriented architecture, comprising:

one or more resources configured to service a business transaction, a plurality of said one or more resources having deployed thereon a protocol stack; and

15 a transaction manager having deployed thereon said protocol stack and being configured to issue in a protocol request header in a protocol request message to said plurality of said one or more resources to monitor one or more performance metrics for said business transaction serviced by said plurality of said one or more resources, each of said plurality of said one or more resources being configured to collect said one or more performance metrics for said business transaction and to aggregate said one or more performance metrics in a protocol response header in a protocol response message that is sent to said transaction manager, wherein said protocol request header in said protocol request message is associated with a thread of execution for said business transaction and wherein said protocol request header in said protocol request message is propagated to any other requests made within said same thread of execution for said business transaction, such that said protocol request header in said protocol request message is replicated through all exchanges associated with said business transaction.

20

25

10. The system according to claim 9, further comprising:

30 a resource manager configured to collect one or more performance metric events from one or more monitoring agents deployed on said one or more resources.

11. The system according to claim 10, wherein said transaction manager is further configured to monitor one or more transaction metrics for said business transaction serviced by said one or more resources and is configured to analyze said one or more transaction metrics monitored for said one or more resources servicing said business transaction in view of said one or more performance metrics collected by said one or more resources.

12. The system according to claim 11, wherein a respective preceding resource of said plurality of said one or more resources is configured to append a respective set of performance metrics to said protocol response header in said protocol response message received from a respective succeeding resource of said one or more resources.

13. The system according to claim 12, wherein said protocol stack corresponding to one or more of said plurality of said one or more resources is pre-configured with policies in order to compare an aggregate value of a performance metric of said one or more performance metrics received in said protocol response header to a pre-set value for said performance metric, and wherein said one or more of said plurality of said one or more resources sends an alert to said resource manager when said pre-set value for said performance metric is exceeded.

14. The system according to claim 13, wherein said protocol stack comprises a Simple Object Access Protocol (SOAP) stack; wherein said protocol request header comprises a Simple Object Access Protocol (SOAP) request header; wherein said protocol response header comprises a Simple Object Access Protocol (SOAP) response header; wherein said protocol request message comprises a Simple Object Access Protocol (SOAP) request message; and wherein said protocol response message comprises a Simple Object Access Protocol (SOAP) response message.

15. A computer program product for monitoring metrics associated with a business transaction, said computer program product comprising:

a computer readable medium;

first program instructions to provide a transaction manager configured to specify, in a request header in a request message issued to one or more resources servicing a business

transaction, collection of one or more performance metrics; each of said one or more resources and said transaction manager having deployed thereon a protocol stack;

second program instructions to collect said one or more performance metrics specified in said request message;

5 third program instructions to send to said transaction manager in a response header in a response message said one or more performance metrics collected by said one or more resources servicing said business transaction; and wherein said first, second and third program instructions are stored on said computer readable medium.

10 16. The computer program product according to claim 15, further comprising:

fourth program instructions to pre-configure said protocol stack deployed on said one or more resources with policies in order to compare an aggregate value of a performance metric of said one or more performance metrics in a respective response header received from a succeeding resource to a pre-set value for said performance metric and to send an alert to a resource manager when said pre-set value for said performance metric is exceeded;

15 said fourth program instructions being stored on said computer readable medium.

17. The computer program product according to claim 16, wherein said second program instructions further include instructions to associate said request header in said request message with a thread of execution for said business transaction and to propagate said request header in said request message to any other request messages made within said thread of execution for said business transaction.

20

18. The computer program product according to claim 17, wherein said second program instructions further include instruction to append a respective set of performance metrics to said response header in said response message received from a succeeding resource of said one or more resources and to send said response message to said transaction manager.

25

19. The computer program product according to claim 18, wherein said transaction manager is further configured to monitor one or more transaction metrics associated with said business transaction and is further configured to analyze said one or more transaction

30

metrics monitored for said business transaction in view of said one or more performance metrics collected from said one or more resources.

5 20. The computer program product according to claim 19, wherein said protocol stack comprises a Simple Object Access Protocol (SOAP) stack; wherein said request header comprises a Simple Object Access Protocol (SOAP) request header; wherein said response header comprises a Simple Object Access Protocol (SOAP) response header; wherein said request message comprises a Simple Object Access Protocol (SOAP) request message; and
10 response message.

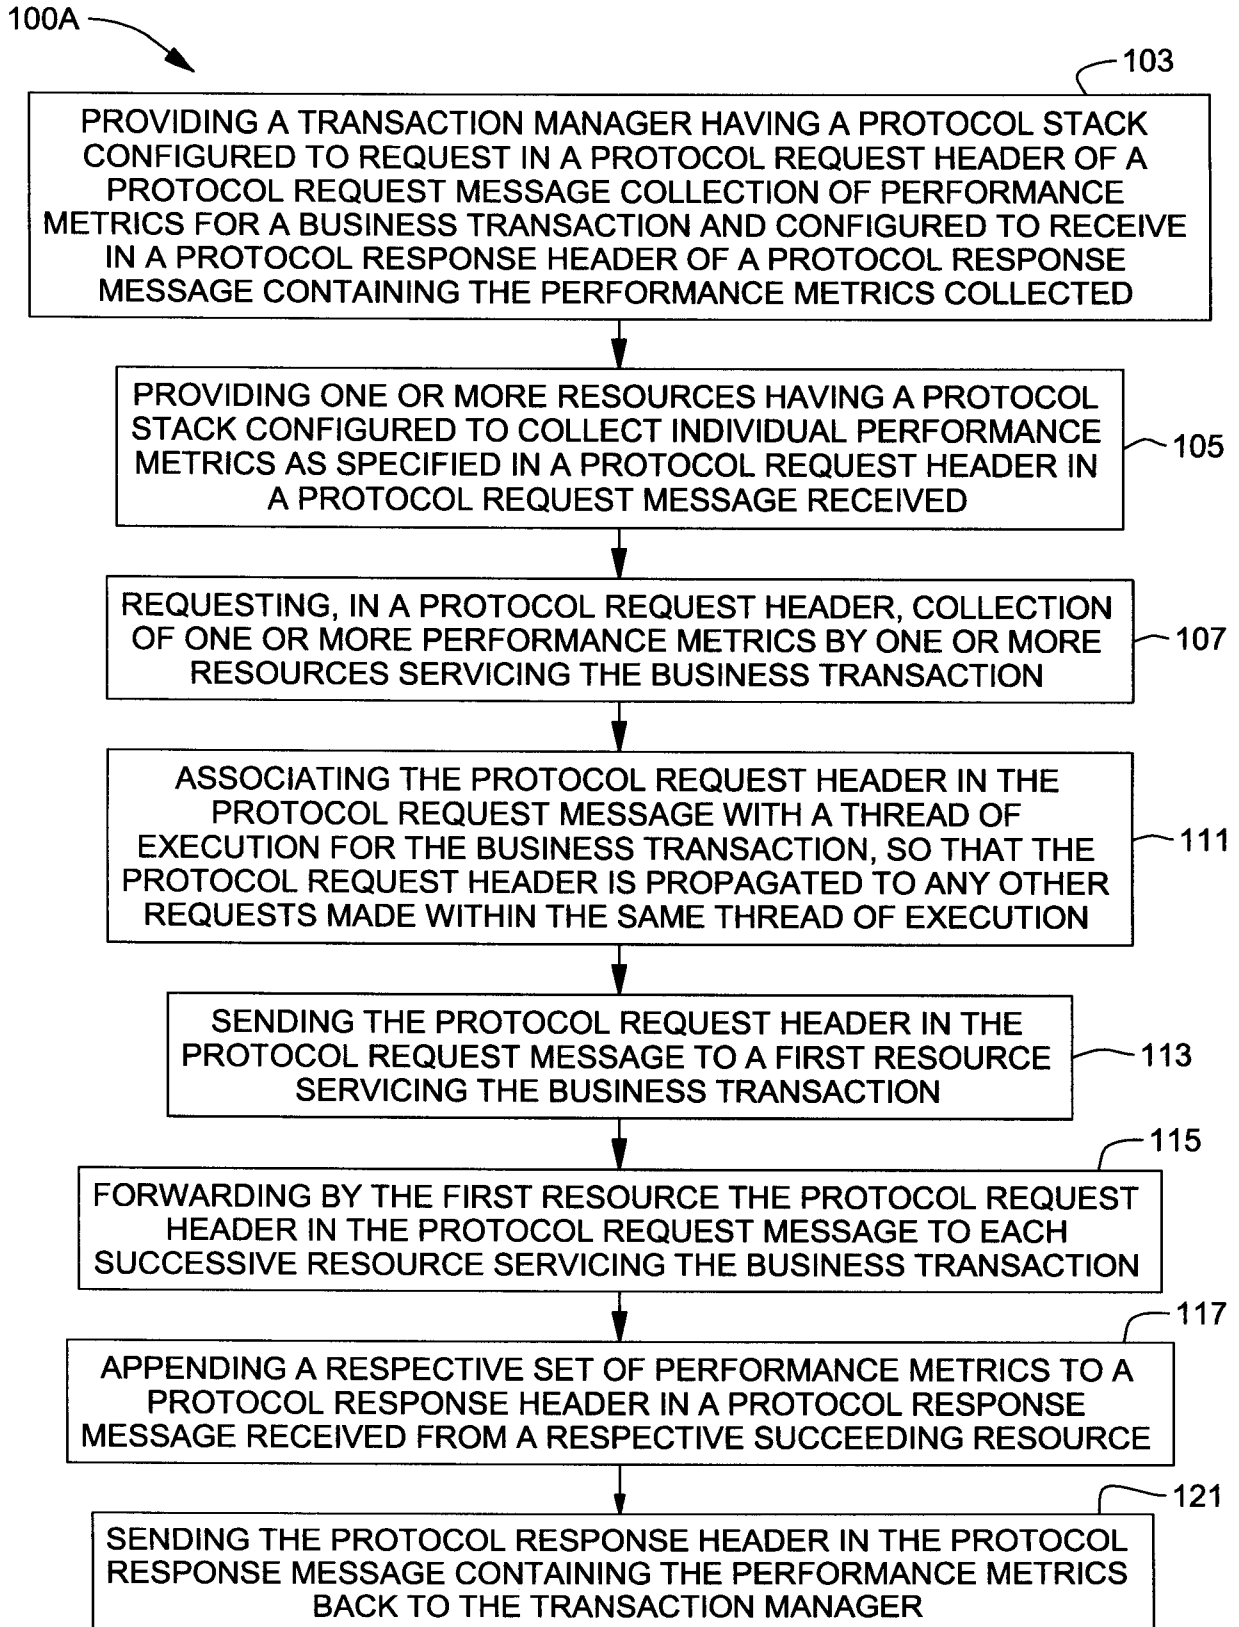


FIG. 1A

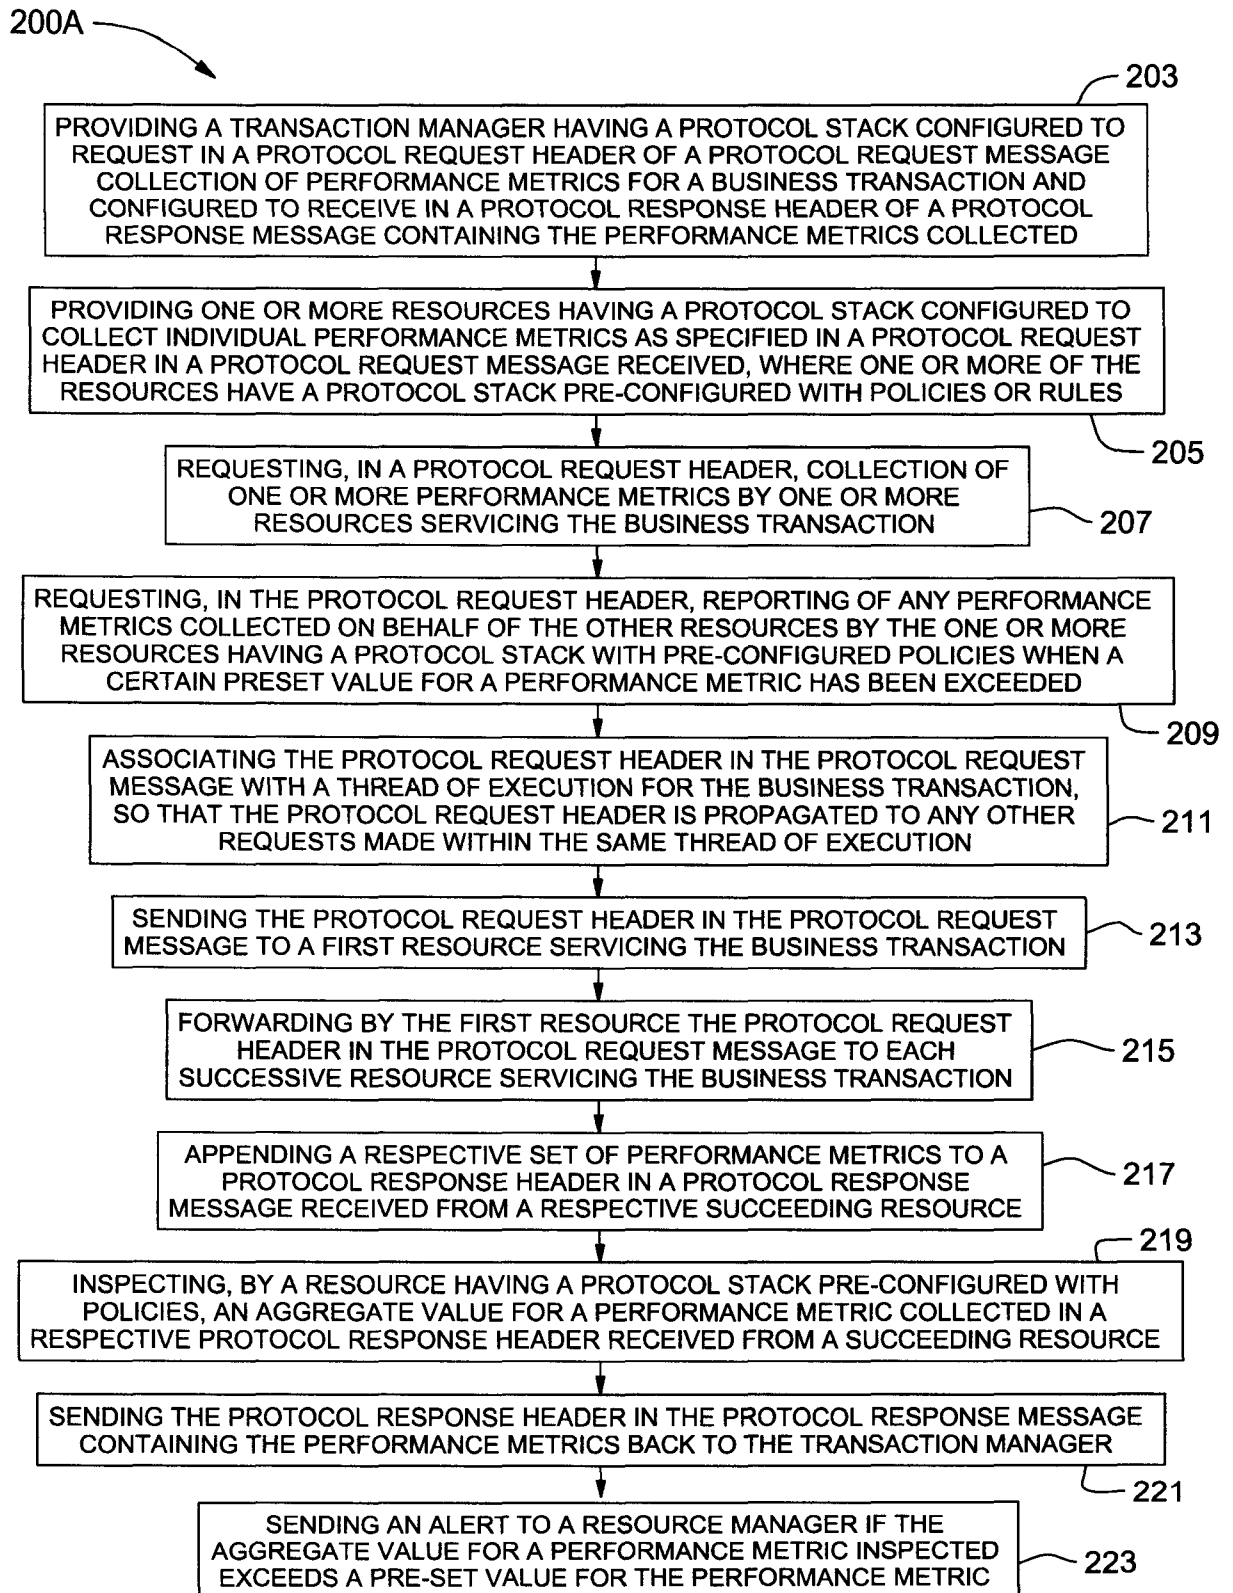


FIG. 2A




FIG. 1B




FIG. 2B




FIG. 3




FIG. 4

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/063538

A. CLASSIFICATION OF SUBJECT MATTER INV. G06F11/34				
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) 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, INSPEC				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
Y	JAN SCHAEFER ED - FRANK ELIASSEN ET AL: "An Approach for Fine-Grained Web Service Performance Monitoring" DISTRIBUTED APPLICATIONS AND INTEROPERABLE SYSTEMS LECTURE NOTES IN COMPUTER SCIENCE;;LNCS, SPRINGER-VERLAG, BE, vol. 4025, May 2006 (2006-05), pages 169-180, XP019034655 ISBN: 3-540-35126-4 abstract; sections 1 and 2 <div style="text-align: center;">----- -/--</div>	1-20		
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.				
<input checked="" type="checkbox"/> See patent family annex.				
* Special categories of cited documents :				
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none; vertical-align: top;"> *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 </td> <td style="width: 50%; border: none; vertical-align: top;"> *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 </td> </tr> </table>			*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
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 <div style="text-align: center;">21 April 2008</div>	Date of mailing of the international search report <div style="text-align: center;">07/05/2008</div>			
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 <div style="text-align: center;">Schneider, Michael</div>			

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2007/063538

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>J A FARRELL ET AL: "Web services management approaches" IBM SYSTEMS JOURNAL, vol. 41, no. 2, 2002, pages 212-227, XP002475061 abstract; page 214, left column, last paragraph; page 214, right column, second paragraph; page 215, left column, lines 3-6; page 215, right again, second paragraph; page 216, right column, last paragraph; page 217, right column; page 219, right column, last paragraph-page 220, left column, second paragraph; page 224, left column-page 225, left column, second paragraph; page 226, right column, first paragraph</p>	1-20
Y	<p>V KALOGERAKI: "Managing distributed objects in peer-to-peer systems" IEEE NETWORK, vol. 1, no. 18, February 2004 (2004-02), pages 22-29, XP002475062 abstract; page 23, left column, last paragraph-page 23, right column, fourth paragraph; page 25, left column, second paragraph-page 26, right column, first paragraph; page 28, right column, last three lines</p>	1-20
Y	<p>US 2005/050549 A1 (JOSEPH JOSHY [US] ET AL) 3 March 2005 (2005-03-03) abstract paragraphs [0018], [0019], [0027]</p>	1-20
Y	<p>US 2005/080930 A1 (JOSEPH JOSHY [US]) 14 April 2005 (2005-04-14) abstract paragraphs [0037], [0041], [0049] - [0051]</p>	1-20
Y	<p>PAPAZOULO M P ET AL: "Web services management: a survey" IEEE INTERNET COMPUTING IEEE USA, vol. 9, no. 6, November 2005 (2005-11), pages 58-64, XP002475063 ISSN: 1089-7801 abstract; page 62, right column, second paragraph</p>	1-20
	<p>----- -/--</p>	

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/063538

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WEI ZHOU ET AL: "Extending SOA to Cover Unreliable Services on Internet" SERVICE-ORIENTED SYSTEM ENGINEERING, 2006. SOSE '06. SECOND IEEE INTERNATIONAL WORKSHOP, IEEE, PI, October 2006 (2006-10), pages 110-113, XP031027456 ISBN: 0-7695-2726-4 abstract; section 4.2	1-20
A	ANDREAS HANEMANN ET AL: "PerfSONAR: A Service Oriented Architecture for Multi-domain Network Monitoring" SERVICE-ORIENTED COMPUTING - ICSOC 2005 LECTURE NOTES IN COMPUTER SCIENCE;;LNCS, SPRINGER-VERLAG, BE, vol. 3826, 2005, pages 241-254, XP019025538 ISBN: 3-540-30817-2 abstract; sections 3 and 4	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2007/063538

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005050549	A1	03-03-2005	NONE
US 2005080930	A1	14-04-2005	CN 1867898 A 22-11-2006
			EP 1683013 A1 26-07-2006
			WO 2005041035 A1 06-05-2005
			JP 2007514990 T 07-06-2007