

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
24 February 2005 (24.02.2005)

PCT

(10) International Publication Number
WO 2005/017715 A2

- (51) International Patent Classification⁷: **G06F**
- (21) International Application Number:
PCT/US2004/026597
- (22) International Filing Date: 16 August 2004 (16.08.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/495,566 15 August 2003 (15.08.2003) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; Intellectual Property Law J46/G468, 555 Bailey Avenue, San Jose, CA 95141 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **LIN-CHO CHAN, Joseph** [GB/US]; 3984 El Nido Ranch Road, Lafayette, CA 94594 (US).
- (74) Agent: **DAVDA, Janaki, Komanduri**; Konrad, Raynes and Victor, 315, S Beverly Drive, Suite 210, Beverly Hills, CA 90212 (US).


- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that 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: METHOD AND SYSTEM FOR MONITORING PERFORMANCE OF PROCESSES ACROSS MULTIPLE ENVIRONMENTS AND SERVERS



(57) Abstract: A method and system in accordance with the invention provides both real time and historical information regarding requests executed on a first server and on a second server. Applications running on two or more servers are monitored. Users are provided with information concerning requests being executed on the two or more servers, which information identifies any request that includes operations performed on more than one server. Information presented to the user includes information identifying two or more servers, processes executed or recently executed on each server in connection with a single request, and methods associated with a single request. Such information may be provided on a single display for a single request, and include information relating to two or more servers. Users may be provided with the ability to drill down to obtain more detailed information.

METHOD AND SYSTEM FOR MONITORING PERFORMANCE OF PROCESSES ACROSS MULTIPLE ENVIRONMENTS AND SERVERS

FIELD OF THE INVENTION

[0001] This invention relates to computer software, and in particular to monitoring performance in connection with distributed computing.

BACKGROUND OF THE INVENTION

[0002] Business-critical applications, including key customer-facing applications, are carried out on servers, i.e., computer programs employing a variety of technologies, including Java J2EE technologies from Sun Microsystems. Methods and systems have been designed and implemented to monitor the behavior of servers for purposes of fault diagnosis and monitoring of efficiency of processing. However, many individual tasks are carried out on more than one server. Indeed, the different servers may employ differing operating systems and application server software. For example, a financial services transaction may involve a Java-based front end and IBM's Customer Information Control System ("CICS") based servers for executing transactions. There are numerous other examples of two and more

[0003] Existing methods and systems can monitor a server based on a single technology, such as a Java-based server. However, those processes and systems cannot monitor and depict for the user an entire task. Existing methods and systems cannot correlate processes carried out on one server with those carried out on another. Thus, it is difficult to identify sources of problems in tasks. For example, if a process appears to be stalled on the Java server, the cause may be in the process itself, or may be because the process is functioning properly, but is awaiting data from the CICS-based server.

SUMMARY OF THE INVENTION

[0004] The invention is a method and system which provides both real time and historical information regarding requests executed on a first server and at least a second server. In the method and system, monitoring of processes includes monitoring of at least one application running on the first server and at least one application running on the second server. The first server and the second server may provide the same environment or a different environment. Users are provided with information concerning requests being executed on at least two of the servers, which information identifies any request that includes operations performed on more than one server. Information as to a single request processed on two or more servers is correlated. Information presented to the user includes information identifying the servers, those operations currently being executed or historically executed on each server in connection with a single request, and methods associated with a single request. Such information may be provided on a single display for a single request, and include information relating to two or more servers. Users may be provided with the ability to drill down to obtain more detailed information.

[0005] The invention is also a system which provides both real time and historical information regarding requests executed on more than one server. In a system of the invention, means are provided for monitoring of at least one application running on a first server and means for monitoring of at least one application running on a second server. The system has means for providing information concerning a request being executed on at least one of said first and second server, and for identifying a request which is executed on both of the first server and the second server. The system has means for correlating information as to a single request processed on two or more servers; means for presenting to the user information identifying the servers, those operations currently being executed or historically executed on each server in connection with a single request, and methods associated with a single request. The system may have means for providing such information on a single display for a single request, and means for displaying information relating to two or more servers. Means may be provided to permit users to drill down to obtain more detailed information.

[0006] The invention is also a computer program for providing information regarding at least one request executed on a first server and a second server. The

program consists of instructions stored on a medium, which instructions, when executed on a processor, cause the processor to execute the steps of monitoring of at least one application running on the first server and at least one application running on said second server, identifying at least one request executed on both the first server and the second servers; and displaying information concerning the identified request.

BRIEF DESCRIPTION OF THE FIGURES

[0007] Fig. 1 is a block diagram illustrating components of the invention applied on a system.

[0008] Figs. 2A and 2B are a flow diagram illustrating a high-level process flow.

[0010] Fig. 3 is a screen shot illustrating a display for a user of all active requests running in a system employing two or more servers.

[0011] Fig. 4 is a screen shot illustrating a display for a user of a composite request screen providing additional details regarding one of the requests listed in Fig. 3.

[0012] Fig. 5 is a screen shot illustrating a display for a user of a composite stack trace display, with additional information regarding portions of a composite request shown in Fig. 4.

[0013] Fig. 6 is a screen shot illustrating a display for a user of a method trace screen listing each method, or equivalent, on each server on which a request shown in Fig. 5 is executed.

[0014] Fig. 7 is a block diagram showing major components in an exemplary architecture for obtaining information regarding requests executed on two servers.

[0015] Fig. 8 is a block diagram illustrating additional features of the exemplary architecture shown in Fig. 7.

[0016] Fig. 9 is a flow diagram illustrating a process flow carried out by a publish server illustrated in Figs. 7 and 8.

[0017] Fig. 10 is a flow diagram illustrating a process flow carried out by a publish server illustrated in Figs. 5 and 6.

[0018] Fig. 11 is a flow diagram illustrating additional details in the process flow of Fig. 10.

[0019] Fig. 12 is a flow diagram illustrating the process of assignment of visa values to a transaction executed by three different servers.

[0020] Fig. 13 is an example of a table containing information regarding transactions executed by two or more servers.

[0021] Fig. 14 is a diagram illustrating the process flow in an exemplary architecture of the invention.

[0022] Fig. 15 is a block diagram illustrating a step in a process flow as carried out in an exemplary architecture.

[0023] Fig. 16 is a block diagram illustrating a step subsequent to that of Fig. 15.

[0024] Fig. 17 is a block diagram illustrating a step subsequent to that of Fig. 16.

[0025] Fig. 18 is a block diagram illustrating a step subsequent to that of Fig. 17.

[0026] Fig. 19 is a block diagram illustrating a final step in the process flow of Figs. 15 to 18.

[0027] Fig. 20 is a block diagram illustrating features of an exemplary architecture in a system in accordance with the invention.

DETAILED DESCRIPTION

[0028] The invention is a method and system which provides both real time and historical information regarding requests executed on more than one server. In the method and system, monitoring of processes includes monitoring of applications running on two or more servers. The method provides a user with information concerning requests being executed on the two or more servers, which information identifies any request that includes operations performed on more than one server. Information as to a single request processed on two or more servers is correlated. Information presented to the user includes information identifying two or more servers, processes executed or recently executed on each server in connection with a single request, and methods associated with a single request. Such information may be provided on a single display for a single task, and include information relating to two or more servers. Users may be provided with the ability to drill down to obtain more detailed information.

[0029] Referring to Fig. 1, the invention will be described as applied to a computer system. The invention is used in connection with a computer system 10 which includes a first server (server 1) 15 which is in communication with a client computer 20. Client computer 20 may be connected via the Internet or other communications network to first server 15. Server 15 is a server having one or more processors, operating system

software, and application software 1 16. System 10 also includes at least second server 2 15', and up to N servers, up to 15 sub N. Second server 15' has application software 2 16', which may be the same as or different from application software 1 15, and server N 15N has application software N 16N. A managing server 30 performs various monitoring and management functions by use of monitoring software 35 running on monitoring server 30. Each of the servers 15 has a data collector program 37, which collects data for monitoring software 35, running on the server. Data collector program 37 provides information regarding applications running on servers 15 to monitoring software 35.

[0030] Managing server 35 has user display 40, which may be a screen display of any suitable type, and user input 42, which may be a keyboard, mouse, or other suitable user input device. Monitoring software 35 displays, on user display 40, information concerning requests being executed on system 10.

[0031] In Fig. 2, a simple flow diagram is provided to illustrate the method of the invention at a high level. In step 50, requests being executed on two or more servers are monitored. This is carried out by the data collectors 37 and monitoring software 35. Each request is assigned an identifier, as indicated at block 52. As requests are passed from one server to another, the identifier is also passed, as indicated at block 53. As requests are passed from one server to another, the fact of the passing of the request, along with the identifier, is provided by the data collector to the managing software, as indicated at block 54. A user, through monitoring software 35, queries the software for a listing of requests. The query is received, as indicated at block 56. The query in this example is for all currently active requests, but the query may be, for example, for historical requests meeting certain parameters, by way of example. The software 35 displays for the user the active requests, identifying those requests that are executed on more than one server, as indicated at block 58. The user may drill down through a series of levels of detail to obtain additional information relating to one of the displayed requests. The user may query the software for additional information regarding a selected one of the displayed requests. As indicated at block 60, such a query is received. The software displays the composite request screen for the request, as indicated at block 62. The user may query the monitoring software for a listing of processes corresponding to the request, as indicated at block 64. The software displays this information, as indicated

at block 66. The user may query the monitoring software 35 for additional detail in the form of a composite method trace. A query for a listing of methods is received, as indicated at block 68. At block 70, the listing of methods is displayed.

[0032] The invention includes a method for providing information relating to the processing of requests, where the processing of requests occurs on more than one server. The multiple servers may include more than one environment. The term “request” is used to include the terms transaction or process, and generally refers to a request made by a client for a one server to take an action. The term “environment” refers to the application software which is processing the request. Examples of environments are WebSphere and CICS. A single request may involve processing taking place on more than one server. It may be that all of the servers are using a single environment, such as technology based on the Java programming language, such as on WebSphere. Alternatively, some of the servers may be a different environment, such as a different application server, such as CICS.

[0033] The method, system and software of the present invention may be used as a component of a larger system, such as that described in published U.S. Patent Applications Publication Nos. US-2004-0054984-A1 and US-2004-0064552-A1. These applications are incorporated by reference herein.

[0034] Software in accordance with the present invention may provide an identifier, such as a numeric identifier, to each request. Preferably, each request is assigned a unique identifier. This identifier is passed by the data collector software when a request is passed between servers. The data collector software also furnishes the identifier to the managing software, so that the managing server can correlate information received from different servers. Requests that move across servers are tracked and identified as such. In the present example, a component of the managing software, called the global publish server, is responsible for maintaining information regarding active requests executed on more than one server.

[0035] The following screen shots are merely examples of the types of displays, and alternative arrangements of information may be employed within the scope and spirit of the invention. It will be appreciated that information regarding all active processes running on more than one server can be shown on a single display. The display may be a

single page provided by a web server or by other suitable display software. Similarly, with respect to each additional level of detail, all of the information may be provided on a single display, which may be a single page provided by a web server or other suitable display software.

[0036] An example of a display of requests, which may be displayed in the step shown at block 58 of Fig. 2A, is shown, for example, at Fig. 3. Fig. 3 shows a display 100, which displays the results of a search sought by a user for all active requests. This may also be called an “in-flight search.” Other criteria may be employed for a search, as indicated by the menus permitting selection of individual servers, individual groups of servers, and particular requests, shown in the search criteria portion 140. The “Search Results” portion 150 of the display identifies certain information regarding all identified active requests. The left hand column 105, labeled “Server Name”, identifies the name of one server on which the request is being executed. In column 105, a designation 110 is provided for requests that are executed in more than one server. As may be seen in this example, a CICS server is included among the servers that are being monitored. Additional information furnished includes an identification 115 of the request, the start date and time 125 of the request, a thread ID 130, and a total resident time 135. The user may drill down into individual listed requests for additional detail, for example by mouse clicking on the thread ID in column 120.

[0037] Referring to Fig. 4, screen display 200 is shown, in which more detailed information than that provided in the display 100 is shown regarding a request being executed on multiple servers. This display is an example of a display provided at the step at block 62 of the process flow of Fig. 4. This screen, called a composite request screen or composite request detail, provides the different names that the request has on the various servers. In this example, there are two such servers. Information provided includes information about the server, such as the server name in the column labeled 220, operating system, in the column labeled 222, and application server, in the column labeled 224, as well as information about the request on that server, such as request name, in the column labeled 230, the request status, in the column labeled 232, the request start time, in the column labeled 234, the request resident time, in the column labeled 236, and the request resident time, in the column labeled 238. The user has the

option of drilling down to obtain additional information about an active portion of a request, such as by clicking on a link, such as at 210.

[0038] Referring to Fig. 5, there is shown a display which may be shown at the step of block 66 of Fig. 2. This display is composite stack trace display 300, with additional detailed information regarding portions of a composite request that are still active. This screen is obtained by user selection from the composite request screen of Fig. 2. The composite stack trace 300 displays information for all of the servers on which the request has executed. For example, those servers may be server 1 15 and server 2 15' of Fig. 1. In the example of Fig. 5, a first server is identified, at 320, as having server name CIQ5, CICS5, and a second server is identified, at 322, as having server name zinx.zenapp-lnx-s02. The operating system and application server program of the first server are also identified at 320. The operating system and application server program of the second server are also identified at 322. The server where the request is active, in this case the server having the server name zinx.zenapp-lnx-s02, lists the processes and their status, as shown at 310. The user can readily see which processes are active and which are inactive.

[0039] The user can drill down and select additional detail from the composite stack trace screen, as indicated in Fig. 2 at block 68, by selecting a composite method trace from the menu 330, which returns a method trace screen, such as the screen 400 illustrated in Fig. 6. The method trace screen 400, which may be displayed in the step at block 70 of Fig. 2, lists each method, or equivalent, on each server on which the request is executed. In this example, the request has been executed on a Java server and a CICS server. Method trace screen lists each method on the Java server, generally at 410, and each program on the CICS server, generally at 430. Information identifying the server is shown, at 412 and 432, respectively. As to each method or program, such information as its elapsed time, CPU time, change in elapsed time, as well as the method or program name is shown. The user can readily see the status of methods and programs in multiple servers.

[0040] It will be understood that the software described below in this application, together with software from published applications nos. US-2004-0054984-A1 and US-2004-0064552-A1, and suitable computer hardware, including processors,

memory, communication devices, and display devices, constitute means for accomplishing the tasks of monitoring certain information and displaying certain information for the user. It will also be understood that the software described below constitutes instructions stored on a medium, which instructions, when executed on a processor, causes the processor to execute the steps necessary to accomplish the tasks described related to monitoring of applications and furnishing and displaying to users certain information.

[0041] An exemplary architecture, in the nature of modifications to the architecture explained in prior applications will be described herein. It will be understood that the architecture and process flow described herein is merely an example of the manner in which the functionality described above may be achieved, and other architectures and process flows may be employed within the scope and spirit of the present invention.

[0042] Referring to Fig. 7, further detail regarding the managing software 35 is shown. Correlation information is maintained by an application called the Global Publish Server 510. This application identifies top level requests, receives information from data collectors, and maintains a database that enables tracking of requests in multiple servers. Fig. 7 is a summary overview, emphasizing the modifications to the architecture disclosed in the above-noted prior application. In this summary example, a method and system of the invention are operating on two servers, namely a WebSphere (R) application server 520, and a CICS application server 525. A data collector 522, configured for the WebSphere (R) application server, and SAM-enabled, as explained below, is collecting information from the WebSphere Application Server. Suitable data collectors are disclosed, for example, in a published U.S. Patent Application under Publication No. US-2004-0054984-A1.

[0043] Data collectors 522, 524 provide information regarding each request to publish servers 540, 542. Publish servers 540, 542, then furnish information to global publish server 510, in addition to furnishing information to other system components. GPS 510 identifies each request with a unique identification number, which may be termed a "passport." A request identification and passport number is then furnished by the GPS to the data collectors. Information provided by the data collectors to the publish

servers then includes the passport number. Information about events is correlated by passport by the GPS. All information as to current requests may be stored as an active GPS itinerary. A recent GPS itinerary may also be stored for recently-completed requests. Historical information is preferably stored in a database. The application server type is maintained in the itinerary. The itinerary information should be sufficiently complete to provide responses to user requests described and illustrated above in connection with the exemplary screen displays of Figs. 3-6.

[0044] As illustrated in Fig. 7, a gateway between the servers 520, 530 is also SAM- enabled; in other words, passport identification information is maintained as requests are transferred across the gateway.

[0045] The GPS 510 receives its information through the publish servers 540, 542, and furnishes the publish servers 540, 542 with the passport number, as well as instructing the publish servers to release the related request method trace when the request is completed. For transactions identified by a passport, a numeric identifier is employed to indicate the number of connections by that transaction between servers tracked by GPS 510, as well as a parent visa number, if any. This numeric identifier is referred to as a "visa," but other terminology may be employed. For example, a first GPS tracked connection may have visa value 0, and a second connection will have visa value 1. If a transaction is called by another transaction with visa value 1, the visa of the first connection made by the called transaction has value 1-0, and the second connection made by the called transaction has value 1-1. In an exemplary embodiment, shown in Fig. 8, a GPS dispatcher is provided as part of the kernel. Publish servers 540, 542, are shown, and a publish server is provided for each monitored server (not shown in Fig. 8). Each publish server can then provide information from the probes to the GPS by updating the GPS dispatcher. In particular, the information from the probes is provided to kernel 1 610. GPS dispatcher 612, running as part of kernel 610, then furnishes the information provided by the publish servers 540, 542, to update each instance of the global publish server. The GPS Dispatcher updates each instance of the GPS. Kernel 2 620 is illustrated, as a second instance of the kernel is preferably running at all times. Second GPS dispatcher 622 is provided as part of Kernel 2. Two instances of the global publish server are also shown as 510 and 510'.

[0046] The data collector will now be discussed. The data collector, in addition to providing functionality such as that described in the above-identified published applications, inserts the GPS passport and GPS visa data into the application data conveyed by the data collector to the publish server. In addition, the data collector provides additional records. The additional records may be the GPS start record, which is sent when the managed server makes a connection to another managed server. The GPS start record contains the passport and visa information, and is furnished to the other managed server. The GPS end record is sent when the managed server disconnects from another managed server. A GPS connect record is sent when a managed server accepts a connection from another managed server. The record may contain passport, visa and component information. Referring to Fig. 9, there is shown a SAM-enabled data collector flow diagram for the sendRequest function in command_agent.C. The process flow starts at 700. For example, if there is a GPS start record, then the data collector will either create a passport, as shown at blocks 708 and 710, or, if there is already a passport, increment a visa, as shown at block 706. In either event, the GPS start data with the passport and visa data is sent to the publish server, as shown at block 712. If there is a GPS connect message, the GPS connect message is simply sent to the publish server, as indicated by blocks 720, 722. If there is an Enterprise Java Bean (EJB) or servlet Start message, the thread data serving request value is incremented, and a request start message will be sent to the publish server, as shown at blocks 730, 732, 734. If there is an EJB/servlet end message, then the thread data serving request value is decreased by 1, as shown at block 744. If the thread data serving request value is then 0, the passport and visa values are set to NULL, and the request end data is sent to the publish server, as indicated at blocks 746, 748. If the thread data serving request value is non-zero, then a request end message is sent to the publish server, without resetting the passport and visa values, as may be seen at block 744 and 748. An example of these functions is described below.

[0047] The publish server will now be described. The publish server is the bridge between the data collectors and the GPS servers. The publish server described in application number US-2004-0054984-A1 may be employed here. There are three record types not used that application, namely GPS start, GPS end and GPS connect. The

processing is illustrated in the Fig. 10 flow diagram. In summary, when a GPS start record is received, if the request is not GPS tracked, then the publish server notifies the GPS (through the kernel) that a cross server application has started. The visa value may be appropriately set to indicate a root transaction.

[0048] When a request end record for the root transaction is received, then the publish server notifies the GPS (through the kernel). The publish server also notifies GPS (through the kernel) when the GPS connect record is received. This record shows the application is connected to another server. The request method trace has a GPS passport assigned and is not discarded.

[0049] Referring specifically to Fig. 10, as shown at block 802, if a GPS start or GPS connect record is received, the next step is to determine if the execution context exists in the mGPS Passport, as shown at block 804. mGPS Passport is a hash map with ExecutionContext as the key and passport String as the value. If so, then processing depends on whether the record is a GPS start or GPS connect record, as shown at block 806. If the record is a GPS start record, then the GPS is called with a visa value of null, as shown at block 808 and the request is added to the mGpsPassport. If the record is a GPS connect record, the request is added to the passport, but the visa value is not set to null, and then the information is forwarded to the GPS. Similarly, if the execution context does not exist in the mGPS Passport, and the record is a GPS connect record, then the GPS is called. For all GPS start records, the processing moves to the Process RLR, at block 816, explained in more detail with respect to Fig. 11. After the ProcessRLR process is complete, if the record is not a GPS record, it is sent to the archive agent, as indicated at blocks 820, 822, and otherwise the process ends.

[0050] In Fig. 11, the handling of method traces by the publish servers, known as processRLR, is set forth in detail. The process commences at block 900, and proceeds to obtaining method traces from the mActiveRequests, as indicated at block 902. If the method trace is NULL, then the request should be a GPS start, and any other result is an error, as indicated by blocks 904, 906 and 908. If the request is a GPS start, then a new method trace is commenced, as indicated at block 910. The request is added to the method trace, as indicated at block 912. If the method trace is not NULL, and the request is a GPS start, as indicated at blocks 904 and 914, then the request is added to the method

trace at block 912. The process then proceeds to push the request to the mRequest StackMap for delivery to the GPS, as indicated at block 916.

[0051] If the method trace is not NULL, and the request is not GPS start or GPS end (i.e., either a GPS connect or an unknown record type), then the process flow ends, as indicated by blocks 920 and 918. For GPS end requests, if the mRequest Stack Map is not empty or the last request is not of an unknown type, as indicated at blocks 920 and 922, then the process flow proceeds to the function pop mRequest StackMap at block 924. If mRequest StackMap is empty, then the request is a request end, and the request number is incremented by 1, as indicated at blocks 926 and 928. If the request is of the same type, the process flow moves to block 924, as indicated at block 926. If the Request Stack Map is not empty, that is the end of processRLR, as indicated at blocks 930 and 932. If the Request Stack Map is empty, then the number of the request is incremented by one, as indicated at block 934. If the execution context does not exist in the GPS Passport, then the request is removed from the active requests, as indicated at blocks 936 and 938. If the execution context does exist in the GPS Passport, then the method trace is put in the GPS Requests, and the GPS is called, before the request is removed from the active requests, as indicated at blocks 940 and 942.

[0052] GPS transactions are sampled as a unit, so that accurate snapshots across servers may be obtained. The sampling rate used is that of the data collector of the first request in the GPS transaction. All related method traces are released and may be archived at the end of an application, so that the method traces may be accurately sampled. If this first request in a GPS transaction is sampled, then the GPS must inform the publish servers of the other requests in the GPS transaction to sample those other requests in the GPS transaction.

[0053] A gateway or interface between the servers carries passport and visa information between the servers. In one example, namely CICS transactions, the CICS Transaction Gateway (CTG) may be employed. The CTG is supported on a variety of operating systems, including MVS, Windows, AIX and Sun.

[0054] For active GPS transactions, the GPS maintains the complete list of transactions in the form of the GPS itinerary. For completed GPS transactions, this information is maintained in the database. In response to a request for a transactional

method trace for a multi-server transaction, the GPS passport and visa for the request are employed to lookup the CORE method trace table. A table mapping between the CORE method trace table entry and the GPS passport/visa is provided.

[0055] A passport includes particular information, which may include component identification, a process identification, thread identification and a time stamp. The component identification and process identification identify a particular server instance, and the thread identification and time stamp uniquely identify a request for that particular server. Thus, the passport identifies the server and request where the high-level request originates. A visa includes a sequence number of the GPS tracked connections. A parent visa number may also be included. Fig. 12 illustrates visa assignment for a request that is processed across three servers. In Fig. 12, the request is started on a WAS (WebSphere (R)) server, indicated generally at 1010, at block 1020. A GPS start record is created, and the visa value is set at zero, as indicated at block 1022. The request is passed to the CICS server 1012, a GPS connect record is created, with the visa value still 0, as indicated at block 1024. The process flow continues to a GPS start record, with the visa value at 0-0, as indicated at block 1026. The request is passed to the IMS server 1014, and a GPS connect record is created, with Visa value of 0-0, as indicated at block 1028. The request is then passed back to the CICS server, and a GPS end record is created, with the visa value at 0-0. The request is then passed back to the IMS server, resulting in a GPS start record being created, and the visa value moving to 0-1, as indicated at 1032. After the request is received at the IMS server, a GPS connect record is created, with the visa value remaining at 0-1, as indicated at block 1034. The request is then passed back to the CICS server, and a GPS end record is created, with the visa value remaining at 0-1. The request is then finally passed back to the WAS server, with a GPS end record being created, and, as the request is on the originating server, the visa value going to 0. It will be understood that the process described in Fig. 12 is merely exemplary.

[0056] An example of a GPS itinerary is shown in Fig. 13, as table 1100. In the passport column 1102, exemplary passport numbers are shown. Each row, such as row 1104, identifies a request. The requests on rows 1104 and 1106 have been processed on the identified WAS server. The request on row 1105 is the same as the request on row

1104, and has been passed to the identified CICS server. The passport number is a key to looking up all transactions associated with a particular request. The other columns are server attributes, which are the CID at column 1108, PID at column 1110, and TID at column 1112 associated with the server, the server type at column 1116, and the visa number associated with the server at column 1118, as well as at least the CID of the publish server associated with the server, at column 1114. The GPS itinerary table may also contain request name, request start time, and CPU clock associated with the request.

[0057] The GPS interfaces include interfaces called when a GPS tracked request is completed so that the GPS can gather information in the request. A publish server notifies the GPS when a new GPS tracked connection is made or when the GPS tracked application is finished. The GPS dispatcher calls the GPS to release all tracked method traces in the publish servers at that time. An array list of all active tracked connections or all recent tracked connections, identified by passport, may be requested from the publish servers. The GPS can then reconstruct the multi-server method trace corresponding to the passport based on method trace from each publish server. The GPS may instruct the publish server to hold the method trace for a particular passport, which is appropriate for an unfinished transaction.

[0058] Referring to Fig. 14, a flow chart of the GPS functionality is provided. In this example, there are two servers; only one instance of the kernel is illustrated. Initially, a new multi-server application is detected running on a first server, by publish server 1 at 1202, and is reported to the kernel 1210 by a GPSAddTrackedRequest. The kernel 1210 forwards the information to all running GPSes, and the GPS itinerary is updated. Publish server 2 1204 reports to the kernel 1210 that the same transaction is tracked on a second server monitored by publish server 2 1204. The kernel 1210 forwards the update to all GPSes. Publish server 1 1202 reports to the kernel 1210 that the root transaction is completed, and the kernel notifies all GPSes 1220, 1222. One GPS, GPS 1220 in this example then calls both publish server 1 1202 and publish server 2 1204 to release all method traces pertaining to the transaction. It will be understood that this process may be carried out with more than two publish servers, and with more than two GPSes.

[0059] A sample workflow will now be explained. It will be understood that this example of a workflow is merely illustrative, and that many other workflows are possible within the scope of the present invention. The components are the same as those illustrated in Fig. 7, and therefore the number from Fig. 7 will be used here. In this example, a servlet updates an account which requires the use of a CICS-based application. The servlet operates on a WebSphere Application server. Referring to Fig. 15, initially, the WebSphere server 520 connects to the CICS server 525 through a CTG. A CORE request start record and a GPS start record are sent by the data collector 522 to the publish server 540 for the WebSphere server 520. The record has passport value p0 and visa value 0. The visa is 0, as this is the first transaction for this passport. The publish server 540 notifies the GPS 510 of the request and the passport number by calling GpsAddTrackedRequest. The visa value is NULL, as this is the root transaction. In this illustration, the kernel is not shown. The thick lines in Fig. 15 show the active components in this step. It will be understood that, in the exemplary architecture, the publish server notifies the kernel, which in turn notifies the GPSes.

[0060] Referring to Fig. 16, in the next step in the process flow of Fig. 15, the CICS server 530 has accepted the connection from the WebSphere server 520. Data collector 532 on the CICS server 530 sends a CORE request start record and a GPS connect record to the publish server 542 associated with the CICS server. The information provided includes the passport number, server information, including CICS CID, PID, and TID, and the publish server CID. The publish server 542 notifies the GPS 510 by providing the passport number, with a visa value of 0. The GPS can correlate the information from the two publish servers using the passport number.

[0061] Referring to Fig. 17, in the next step in the process flow, the portion of the transaction executed on the CICS server is now completed. The data collector 532 on the CICS server reports to the associated publish server 542 using a CORE request end record, which includes passport information. The publish server 542 notifies GPS 510 that the request is completed on the CICS server. In particular, the publish server 542 calls GpsTrackedRequestFinished to notify GPS 510 that the request is completed. The GPS updates its itinerary with relevant data, such as CPU time used for the request, and elapsed time for the request, as to the CICS server.

[0062] Referring to Fig. 18, in the next step of the process flow, the connection between the WebSphere server 520 and the CICS server 530 as to this transaction or request is illustrated. The data collector 522 on the WebSphere server 520 sends a GPS request end record, including passport information to the publish server 540 associated with the WebSphere server. The publish server does not need to notify the GPS 510 of the disconnection until the request has been completed. The GPS start and end records are maintained in the method trace for reconstruction of the transactional method trace.

[0063] Referring to Fig. 19, in the final step of the illustrated process flow, the underlying transaction has been completed. The data collector 522 on the WebSphere server sends a CORE request end record to the associated publish server 540. The publish server 540 notifies GPS 510 that the transaction is completed. Publish server 540 may send this information by calling `GpsTrackedApplicationFinished`. The GPS 510 then transfers all records with the associated passport number to the Recent GPS itinerary, and calls both publish servers to release the related method traces.

[0064] Additional features of an exemplary architecture that may be used with the method of the invention may be understood with reference to Fig. 20. In Fig. 20, the monitored servers and the kernel are not shown. It will be understood that Fig. 20 is an example of requests performed on both a Java server and a CICS server, although the invention is not limited to two servers or to Java and CICS. In particular, it can be seen from Fig. 20 that each publish server 540, 542, archives event data to an associated archive agent 550, 552. This occurs at the request sampling rate for data provided by the initiating request server, but either at the request sampling rate or as supplemental data to the initial request when obtained from the publish server of the second server. The archive agents then provide their data to the database 560.

[0065] Continuing to refer to Fig. 20, the visualization engine 570 furnishes information in response to a user request. When rendering composite request detail pages and composite method trace pages for active composite requests, the visualization engine first obtains itinerary information from the GPS. Based on the itinerary information, the visualization engine 570 then identifies the publish servers that maintain the required information about the active composite requests, and then obtains the information as to active composite requests from each publish server 540, 542. However, if the composite

request is completed, the visualization engine 570 obtains the information from the database 560. In the database, only the initial request may be identified as having been part of a composite request. However, the other requests can be obtained as they are correlated with the initial request in the database.

[0066] Referring further to Fig. 20, exemplary communications protocols that may be employed are illustrated. For example, the data collectors 522, 532 employ TCP/IP socket protocols to furnish information to publish servers 540, 542. Data collectors 522, 532 publish GPS data to publish servers as part of event data packets which the data collectors customarily furnish to the publish servers. The data collector on the Java server connects to the CTG through a TCP/IP socket and through an HTTP protocol. The CTG routes the request to the CICS region using an HTTP protocol and EXCI. On a single server, a participating request can invoke another request through DPL, using a local method invocation. Publish servers 540, 542, maintain internal caches of active requests using local method invocation. The publish servers 540, 542 forward GPS data to global publish server 510 using RMI. Global publish server 510 maintains active GPS itineraries internally using local method invocation. GPS 510 sends instructions to publish servers to archive or release participating requests when the home request ends using RMI. Publish servers 540, 542 send home request and participating request event data at the request sampling rate to archive agents 540, 542, using a TCP/IP socket protocol. Publish server 542, which in this example is associated with a server which is not the request's home server, can send, to archive agent 552, the participating request's event data as supplemental data to any sampled home request. Visualization engine 570 uses RMI to get active home request's event data from publish server 540, to get active participating requests' event data from publish server 542, and to get the GPS itinerary from GPS 510. Archive agent 550 uses JDBC to send home request event data to database 560. Archive agent 552 also uses JDBC to send participating request event data to database 560. Visualization engine 570 queries database 560 for completed home request and participating request event data using JDBC.

[0067] The method, system and computer program of the invention enables a user to track a transaction through from start to finish, even when the transaction involves multiple servers. The user can better diagnose sources of problems in execution of

transactions, and can more readily analyze requirements, such as needs for additional capacity.

[0068] The present invention can be embodied in the form of methods and apparatus for practicing those methods. The present invention can also be embodied in the form of program code embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other machine-readable storage medium, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. The present invention can also be embodied in the form of program code, for example, whether stored in a storage medium, loaded into and/or executed by a machine, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. When implemented on a general-purpose processor, the program code segments combine with the processor to provide a unique device that operates analogously to specific logic circuits.

[0069] While the exemplary embodiments have been described with respect to enterprise applications using Java technologies, and specifically J2EE technologies, the methods of the invention may be implemented in other technologies. For example, the methods of the invention may be implemented in other platform-independent technologies. The methods of the invention may also be implemented in other technologies applicable to distributed enterprise applications.

[0009] While the invention has been described with reference to preferred embodiments, the invention should not be regarded as limited to preferred embodiments.

What is claimed is:

1. A method for providing information regarding at least one request executed on a first server and a second server, comprising:
 - monitoring of a first application running on said first server and a second application running on said second server;
 - identifying at least one request executed on both said first server and said second servers; and
 - displaying information concerning said identified request.
2. The method of claim 1, wherein said step of displaying comprises displaying said information on a single display.
3. The method of claim 1, further comprising displaying, in response to a user query, at least one process on said first server associated with said identified request and at least one process on said second server associated with said identified request.
4. The method of claim 3, further comprising displaying, in response to a user query, at least one method corresponding to said identified request.
5. The method of claim 1, further comprising associating an identifier with said request.
6. The method of claim 5, wherein, when said request is passed between said first server and said second server, said identifier is passed with said request.
7. The method of claim 6, wherein, when said request is passed between said first server and said second server, said identifier and associated information is furnished by a data collector to a managing server.
8. A system for providing information regarding at least one request executed on a first server and a second server, comprising:
 - means for monitoring of a first application running on said first server and a second application running on said second server;
 - means for identifying at least one request executed on both said first server and said second server; and
 - means for displaying information concerning said identified request.
9. The system of claim 8, wherein said means for displaying comprises means for displaying said information on a single display.

10. The system of claim 8, further comprising means for displaying, in response to a user query, at least one process on said first server associated with said identified request and at least one process on said second server associated with said identified request.
11. The system of claim 10, further comprising means for displaying, in response to a user query, at least one method corresponding to said identified request.
12. The system of claim 8, comprising means for associating an identifier with said request.
13. The system of claim 12, comprising means for, when said request is passed between said first server and said second server, passing said identifier with said request.
14. The system of claim 13, comprising means for, when said request is passed between said first server and said second server, furnishing said identifier and associated information to a managing server.
15. A computer program for providing information regarding at least one request executed on a first server and a second server, said program consisting of instructions stored on a medium, said instructions, when executed on a processor causing the processor to execute the steps of:
 - monitoring of at least one application running on said first server and at least one application running on said second server,
 - identifying at least one request executed on both said first server and said second servers; and
 - displaying information concerning said identified request.





FIG. 1


FIG. 2A



2/21

FROM FIG. 2A

FIG. 2B



Publishing Composite Request Data: Component Interactions




FIG. 20

4/21

In-Flight Request Search - Multi-Thread Explorer

URL: <http://www.buratt.com/active>

Cyanea

SEARCH CRITERIA

Start: 2002-01-01 00:00:00
End: 2002-01-01 00:00:00
Search Request/Response

SEARCH RESULT

Time/Date	Start/End Time	Thread ID	Total Resident Time (ms)
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	1316100340	6700
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	1316100340	6702
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	1316100340	6682
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	1316100340	6680
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	1316100340	5157
2002-01-01 00:00:00 - 2002-01-01 00:00:00	2002-01-01 00:00:00 PM	00	4900

Serial No.: To Be Assigned
In re: Joseph Lim-Cho Chan

140

150

105

110

115


F16.3

125

130


135


12






F16.6





13

Fig 9



handleRequestRecord






FIG. 10



Visa Assignment



Serial No.: To Be Assigned
In re: Joseph Lim-Cho Chan
Sheet

1102 ✓ 1114 ✓ 1116 ✓ 1118 ✓


Passport	CID	FID	TID	PSCID	Type	Visa
12345	24	15	23	12	WAS	0
12345	25	13	20	13	CICS	1
32350	24	15	24	12	WAS	0

FIG. 13


110

110

110




Serial No.: To Be Assigned
In re: Joseph Lim-Cho Chan
Sheet



32


532

2




Serial No.: To Be Assigned
re: Joseph Lim-Choo Chan

Serial No.: To Be Assigned
In re: Joseph Lim-Cho Chan
Sheet



332



Serial No.: To Be Assigned
In re: Joseph Lim-Cho Chan

Sheet

