



(22) **Date de dépôt/Filing Date:** 2005/04/07

(41) **Mise à la disp. pub./Open to Public Insp.:** 2005/12/04

(45) **Date de délivrance/Issue Date:** 2015/12/01

(30) **Priorité/Priority:** 2004/06/04 (US60/576,805)

(51) **Cl.Int./Int.Cl.** *H04L 12/26* (2006.01),
G06F 11/30 (2006.01), *G06F 11/34* (2006.01),
H04L 12/24 (2006.01)


(72) **Inventeurs/Inventors:**
ALON, AMIR, IL;
LAVI, YORAM YORI, US;
TAL, MORDECHAI, IL

(73) **Propriétaire/Owner:**
SAP SE, DE

(74) **Agent:** SMART & BIGGAR

(54) **Titre : SYSTEME ET METHODE DE GESTION DE PERFORMANCE DANS UN ENVIRONNEMENT INFORMATIQUE MULTICOUCHE**

(54) **Title: SYSTEM AND METHOD FOR PERFORMANCE MANAGEMENT IN A MULTI-TIER COMPUTING ENVIRONMENT**



(57) **Abstr  /Abstract:**

Apparatus for monitoring a selected tier in a multi-tier computing environment, the apparatus including a context agent and a dynamic tier extension, the context agent associated with the selected tier and coupled with other context agents, each of which is associated with a respective tier, the dynamic tier extension coupled with the context agent and with specific predetermined points of the selected tier, the dynamic tier extension monitoring request traffic passing through the selected tier, the monitored request traffic including at least one entering request received at a request entry port from an adjacent tier, the dynamic tier extension identifying each request in the monitored request traffic and sending at least a request identifier to the context agent, the context agent also receiving information relating to the request context of the entering request from the context agent associated with the adjacent tier, the context agent associating the information relating to the request context of the entering request with the entering request, in accordance with the received request identifier.



ABSTRACT

Apparatus for monitoring a selected tier in a multi-tier computing environment, the apparatus including a context agent and a dynamic tier extension, the context agent associated with the selected tier and coupled
5 with other context agents, each of which is associated with a respective tier, the dynamic tier extension coupled with the context agent and with specific predetermined points of the selected tier, the dynamic tier extension monitoring request traffic passing through the selected tier, the monitored request traffic including at least one entering request received
10 at a request entry port from an adjacent tier, the dynamic tier extension identifying each request in the monitored request traffic and sending at least a request identifier to the context agent, the context agent also receiving information relating to the request context of the entering request from the context agent associated with the adjacent tier, the context agent
15 associating the information relating to the request context of the entering request with the entering request, in accordance with the received request identifier.

77964-10

SYSTEM AND METHOD FOR PERFORMANCE MANAGEMENT IN A MULTI-TIER COMPUTING ENVIRONMENT

FIELD OF THE DISCLOSED TECHNIQUE

The disclosed technique relates to managing distributed computing environments in general, and to a system and method for application performance management in a multi-tier computing environment, in particular.

BACKGROUND OF THE DISCLOSED TECHNIQUE

Information Technology (IT) is responsible for delivering application services using an increasingly complex multi-tier production environment with a heterogeneous application mix. IT operations are struggling to meet required service levels in performance and availability, while being pressed to increase efficiency and resource utilization. Consolidation of IT resources, together with business concerns, exacerbates this effect, stretching the capability of IT operations to meet ever changing demands for computing resources. Traditional approaches

and tools for performance and availability management are variations of the never-ending "monitor-tune-fix" cycle, which involves identifying that a problem exists (i.e., monitoring), increasing overall throughput to overcome the problem (i.e., tuning), and performing root-cause analysis to uncover the precise cause of each specific instance of a problem (i.e., fixing). Such approaches are unable to cope with the complexity and variability of the rapidly changing IT environment.

Reference is now made to Figure 1, which is a schematic illustration of a multi-tier computing environment, generally referenced 50, which is known in the art. Computing environment 50 includes a first client 62 running a first application, a second client 64 running a second application, a first tier 52, a second tier 54, a third tier 56, a fourth tier 58, and a fifth tier 60. The first tier 52 is a web server. The second tier 54 is an application server, application server A. The third tier 56 is another application server, application server B. The fourth tier 58 is a further application server, application server C. The fifth tier 60 is a database. First tier 52 is coupled with first client 62, with second client 64, and with second tier 54. Second tier 54 is further coupled with third tier 56 and with fourth tier 58. Third tier 56 is further coupled with fourth tier 58 and with fifth tier 60. Fourth tier 58 is further coupled with fifth tier 60.

A "tier" represents a certain type of processing that is part of the overall delivery of an IT service (e.g., presentation level processing on a

002412CA

web server tier or data processing on a database tier). Each tier typically runs on a different host machine.

The first application initiates a user request R1 and sends user request R1 to first tier 52. User request R1 is part of an overall transaction initiated by the user. User request R1 may be, for example, a web based query to retrieve information from a certain application. User request R1 may require the services of different tiers in computing environment 50 and may generate additional requests in order to obtain these services. The tier that receives a request either replies to the tier that sent the request, or sends a new request to a different tier. Eventually a reply is returned in response to the original user request R1. A given tier can only request a service from another tier in computing environment 50 if the two tiers are directly coupled with each other.

Overall management of distributed computing environment 50 requires knowledge of how each tier handles its workload. For example, given a shortage of resources on one tier, a system administrator may scale this tier by creating clones of the tier, both vertically (i.e., within the same host machine) and horizontally (i.e., across multiple host machines). For example, in computing environment 50, the system administrator may add an additional application server A2 (not shown) to second tier 54 application server A, wherein application server A2 is a clone of application server A. By the same token, if an overabundance of resources exists on a tier, the system administrator may transfer free

002412CA

resources to another tier which has a shortage of resources. The system administrator may further configure a certain tier in order to improve the overall performance or indicate modifications to optimize the application running on the tier. This is an example of tier specific application
5 monitoring for performance management. It is noted that a request might reach only certain tiers in computing environment 50. Furthermore, the same request might reach certain tiers using multiple paths. For example, in computing environment 50, a request may reach fifth tier 60 database via either third tier 56 application server B, or via fourth tier 58 application
10 server C. As the request paths are not consistent across the entire environment, solving the resource shortage on one tier does not necessarily guarantee the performance of the overall application, which may span multiple tiers. A processing bottleneck in any tier will delay all application functions that depend on that tier.

15 First tier 52 receives user request R1. First tier 52 allocates processing enclave X1 to process user request R1. While processing user request R1, the application logic executing in processing enclave X1 determines it cannot complete processing user request R1 without additional information or operations to be provided by second tier 54. First
20 tier 52 then sends a subsequent request R2 to second tier 54, requesting the additional information or operations. Second tier 54 allocates processing enclave X2 to process request R2. The application logic executing in processing enclave X2 determines that request R2 requires

002412CA

further information or operations to be provided by fourth tier 58. Second tier 54 then sends a subsequent request R3 to fourth tier 58. Fourth tier 58 allocates processing enclave X4 to process request R3.

Processing enclave X4 completes execution. Fourth tier 58
5 returns a reply R3' to second tier 54, in response to earlier request R3 of second tier 54. Processing enclave X2 receives reply R3' and resumes processing. Once processing enclave X2 has completed execution, second tier 54 returns a reply R2' to first tier 52, in response to earlier request R2 of first tier 52. Processing enclave X1 receives reply R2' and
10 resumes processing. Once processing enclave X1 has completed execution, first tier returns a reply R1' to user request R1, whose service has now been completed.

In computing environment 50, each of the different tiers is isolated from the tiers which are not directly coupled therewith. For
15 example, request R3 from second tier 54 to fourth tier 58, directly coupled therewith, does not necessarily include information relating to a former request R2, which was received in second tier 54 from first tier 52, nor does request R3 include information relating to user request R1. A given tier has no way of obtaining certain information related to the request
20 being processed at that tier, such as which user initiated the transaction, which requests preceded the request which is being processed at the given tier, or characteristics of requests which preceded that request. For example, second tier 54 cannot identify characteristics of request R2, such

as whether the request was preceded by user request R1 sent to first tier 52, or that the transaction originated at user request R1 from first application 62. As a result, if a priority level is assigned to a processing enclave processing a request within a certain tier, that priority level is assigned taking into account only the minimal information which is available on the tier. This information includes the request characteristics (e.g., the tier login credentials used by the request) and perhaps information about the processing enclave processing that request (e.g., the database session identification). Requests are generally processed on an equal priority basis (e.g., first-come-first-serve), though mechanisms operating to differentiate priority levels are available locally on a given tier. Performance management must be done on an individual tier basis, as the other tiers in computing environment 50 cannot be accounted for when dealing with a specific tier. Typically, a system administrator who is responsible for managing a multi-tier computing environment such as computing environment 50 attempts to improve performance by adjusting the resource allocation for a given tier.

US Patent No. 5,958,010 to Agarwal et al. entitled "Systems and methods for monitoring distributed applications including an interface running in an operating system kernel", is directed to systems and methods for monitoring enterprise wide operation of a distributed computing system to develop business transaction level management data for system performance, usage trends, security auditing, capacity

planning, and exceptions. A system having a distributed computing architecture includes multiple workstations, servers, and network devices. Each workstation is representative of a computer system coupled to a network. Each workstation is capable of requesting service from any of
5 the servers. Each workstation has a communication stack for exchanging data with the network. The system further includes a plurality of monitoring agents, and a console module with a database connected therewith. Each monitoring agent has an external event interface that provides event information about various components of an enterprise.
10 Each of the monitoring agents is associated with a respective one of the workstations or servers.

The monitoring agent may physically reside on the associated client or server thereof. The monitoring agent monitors and collects data being exchanged between a client and the network, and between a server
15 and the network. Each monitoring agent can be a software module, a hardware device, or a combination thereof. Each monitoring agent passes information representative of the collected data to the console module. The console module stores this information within the database for analysis by an operator. An application program running on the console
20 module can view the collected data to show system performance of any process or component of the enterprise. A system administrator can develop enterprise level usage statistics and response times, develop charts and reports, and perform other relevant data analysis for

determining user-defined statistics relevant to the operation of the enterprise.

US Patent No. 6,108,700 to Maccabee et al entitled "Application end-to-end response time measurement and decomposition", is directed to a method and system for measuring and reporting availability and performance of end-to-end business transactions. The system operates on a client-server application architecture. The system includes three logical components: Event Generation, Transaction Generation, and Report Generation, as well as overall system management via System Administration.

The Event Generation component exists on every computer being measured in the architecture. Each computer has one Agent, a plurality of Sensors and a plurality of Processors. The Sensors interact with platform components on which business applications run, monitor application activities, and detect changes of state. When appropriate, each of the Sensors generates an event that describes the change in state, when and where the event occurred, and any extra data necessary to uniquely identify the event. An event contains a time-stamp and correlation data used later by the system to associate the event with other events into transactions. The Sensors forward the generated events to their respective Agents. The Agents temporarily store the data and may distribute the data to other system components having registered interest in the event. A Processor analyzes the events and further deduces

changes in state. The changes in state may be directly related to actions occurring within the business transaction platform components or derived by combining previously generated events from Sensors or other Processors to describe states achieved. The Processors forward the
5 generated events to their respective Agents.

The Transaction Generation component typically exists in one of the computers in the network and includes a Director. The Director receives events from the Agents under control thereof. The events are examined, and correlated and collated into transactions based on
10 transaction generation rules. The System Administrator determines which transactions to generate.

The Report Generation component includes a Manager. The Manager collects the transactions from the Directors. The collected transactions are manipulated to obtain information relating to the
15 availability and performance of business transactions. A report or continuous graphic monitoring can be produced upon a specific or periodic request from a Graphical User Interface (GUI). Report Generation includes definition of the initial selection and processing of transactions, as well as the sorting and aggregation methods used to consolidate the
20 transactions event data into availability and performance information.

US Patent Application No. 2002/0129137 A1 to Mills et al. entitled "Method and system for embedding correlated performance measurements for distributed application performance decomposition", is

002412CA

directed to techniques for embedding correlated performance measurements in transactions associated with a distributed application. The techniques are used in accordance with application performance decomposition. Data is embedded in a communications protocol used to

5 carry a transaction between application components in a distributed computing network, rather than altering the actual transaction data itself. The embedded data may include a timestamp and duration measurement data. The format of the embedded data combines a well-defined keyword prefix with a variable suffix that identifies the timing source, followed by a

10 colon delimiter and whitespace, and followed by the time stamp and duration information.

Subsequent processing stages of the distributed application can interpret the communications protocol to glean processing durations of previous stages, in order to make decisions regarding treatment of the

15 transaction. The measurement information is embedded within the same distributed application described by the measurement information, so that completion of the transaction occurs simultaneous or contemporaneous with availability of knowledge of the transaction performance characteristics.

20 A possible communications protocol is the HyperText Transport Protocol (HTTP). A possible distributed computing network is the World Wide Web (WWW). The application components may be a client application running on a client and a server application running on an

002412CA

application server. For example, the client application is a web browser, and the server application runs on a web server. An application transaction is the client application requesting content from the application server and the application server responding. Performance information is
5 generated to measure the round trip response time from the perspective of the client application, as well as to decompose the response time into the time taken by the server application to service the request and generate a reply. In particular, lines are added to the HTTP headers to carry performance measurement data, allowing the client to receive the server
10 measurement duration in the HTTP Reply header.

SUMMARY OF THE DISCLOSED TECHNIQUE

In accordance with the disclosed technique, there is thus provided an apparatus for monitoring a selected tier in a multi-tier computing environment. The apparatus includes a context agent and a
5 dynamic tier extension. The context agent is associated with a selected tier, and is coupled with other context agents, each of which is associated with a respective tier. The dynamic tier extension is coupled with the context agent and with specific predetermined points of the selected tier, such as at least a request entry port of the selected tier. The dynamic tier
10 extension monitors request traffic passing through the selected tier, the monitored request traffic including at least one entering request received at a request entry port from an adjacent tier. The dynamic tier extension identifies each request in the monitored request traffic and sends at least the request identifier to the context agent. The context agent also receives
15 information relating to the request context of the entering request from the context agent associated with the adjacent tier. The context agent associates the information relating to the request context of the entering request with the entering request, in accordance with the received request identifier.

20 The dynamic tier extension may further be coupled with a request exit port of the selected tier, and the monitored request traffic further includes at least one exiting request exiting the selected tier to an adjacent tier. The dynamic tier extension captures the request context of

the exiting request and sends at least the request identifier and information relating to the request context of the exiting request to the context agent. The context agent associates the entering request with the exiting request, in accordance with the received information and request
5 identifiers. The context agent further provides information relating to the request context of the exiting request to the context agent associated with the adjacent tier to which the exiting request was sent.

In accordance with the disclosed technique, there is further provided a system for application performance management in a multi-tier
10 computing environment including a plurality of tiers. The system includes, for each of at least two monitored tiers of the plurality of tiers, a respective dynamic tier extension and a respective context agent. The dynamic tier extension is coupled with at least a request entry port of the monitored tier. The context agent is coupled with the dynamic tier extension and with
15 other context agents associated with the tiers that are directly coupled with the monitored tier. The dynamic tier extension monitors request traffic passing through the selected tier, the monitored request traffic including at least one entering request received at a request entry port from an adjacent tier. The dynamic tier extension identifies each request in the
20 monitored request traffic and sends at least the request identifier to the context agent. The context agent also receives information relating to the request context of the entering request from the context agent associated with the adjacent tier. The context agent associates the information

002412CA

relating to the request context of the entering request with the entering request, in accordance with the received request identifier. The system further includes a context network management server. The context network management server is coupled with the context agents. The context network management server collects and analyzes performance data received from the context agents.

In accordance with the disclosed technique, there is further provided a method for application performance management in a multi-tier computing environment including a plurality of tiers. The method includes, for each of at least two monitored tiers of the plurality of tiers, the procedure of receiving information relating to the request context of at least one entering request, the information including at least a request identifier and a transaction identifier. The method further includes the procedure of monitoring request traffic passing through the monitored tier, the monitored request traffic including at least the entering request. The method further includes the procedures of identifying the entering request in accordance with the request identifier, and associating the entering request with a transaction in accordance with the transaction identifier.

In accordance with the disclosed technique, there is further provided another method for application performance management in a multi-tier computing environment including a plurality of tiers. The method includes, for each of at least two monitored tiers of the plurality of tiers, the procedure of monitoring request traffic passing through the monitored tier,

77964-10

the monitored request traffic including at least an entering request and an exiting request, the exiting request sent from the monitored tier to an adjacent tier. The method further includes the procedures of determining information relating to the request context of the entering request, and identifying each request in the monitored
5 request traffic. The method further includes the procedures of associating the entering request with the exiting request, and sending information relating to the request context of the exiting request to a context agent associated with the adjacent tier.

According to one aspect, there is provided an apparatus for monitoring
10 a selected tier in a multi-tier computing environment, said selected tier is coupled with at least one adjacent tier of said multi-tier computing environment, said apparatus comprising: a context agent associated with said selected tier of said computing environment, said context agent being coupled with other context agents, each of said other context agents associated with a respective one of said at least one
15 adjacent tier; and a dynamic tier extension, said dynamic tier extension coupled with said context agent and with at least a request entry port of said selected tier, at least for monitoring request traffic passing through said selected tier, said monitored request traffic including at least one entering request received at said request entry port from a sending tier of said at least one adjacent tier, said dynamic tier extension
20 providing request information to said context agent, said request information enabling said context agent to identify each request included in said monitored request traffic, wherein said context agent receives context information directly from another context agent associated with said sending tier of said at least one adjacent tier, said received context information relating to the request context of said at least one
25 entering request, said received context information including at least a unique transaction identifier assigned to each request of a transaction to which said at least one entering request belongs; and wherein for each said at least one entering request, said context agent associates said context information with said request information, respectively.

77964-10

According to another aspect, there is provided an apparatus for monitoring a selected tier in a multi-tier computing environment, said selected tier is coupled with at least one adjacent tier of said multi-tier computing environment, said apparatus comprising: a context agent associated with said selected tier, said context agent being coupled with at least one other context agent, each of said at least one other context agent being associated with a respective one of said at least one adjacent tier of said computing environment; and a dynamic tier extension, said dynamic tier extension coupled with said context agent and with at least a request entry port and a request exit port of said selected tier, said dynamic tier extension at least monitoring request traffic passing through said selected tier, said monitored request traffic including at least one entering request received at said request entry port and at least one respective exiting request at said request exit port, said selected tier providing said at least one respective exiting request to a receiving tier of said at least one adjacent tier, said dynamic tier extension capturing a request context of said at least one entering request, said dynamic tier extension assigning a unique transaction identifier to said at least one entering request for uniquely identifying each request of a transaction to which said at least one entering request belongs, said dynamic tier extension providing request information and context information to said context agent, said request information at least identifying each request included in said monitored request traffic, said context information relating to said captured request context of said at least one entering request, said context information including at least said unique transaction identifier, wherein said context agent associates said at least one entering request with said at least one exiting request, and wherein said context agent provides context information relating to the request context of said at least one respective exiting request directly to another context agent associated with said receiving tier of said at least one adjacent tier.

According to still another aspect, there is provided in a multi-tier computing environment including a plurality of tiers, a system for application performance management, associated with at least two selected monitored tiers of said plurality of tiers, the system comprising: for each of at least two of said

77964-10

monitored tiers, a respective dynamic tier extension being coupled with at least a request entry port of said monitored tier, said dynamic tier extension at least monitoring request traffic passing through said respective monitored tier, said monitored request traffic including at least one entering request received at said request entry port; and for each of said at least two of said monitored tiers, a

5 respective context agent, coupled with said respective dynamic tier extension, said context agent being coupled with other said context agents associated with tiers that are directly coupled with said respective tier, said respective context agent receiving request information from said respective dynamic tier extension, said request

10 information at least identifying each request included in said monitored request traffic, said respective context agent further receives context information relating to the request context of said at least one entering request directly from another context agent associated with a respective one of said tiers that are directly coupled with said respective tier, said received context information including at least a unique

15 transaction identifier for uniquely identifying each request of a transaction to which said at least one entering request belongs, for each said at least one entering request, said respective context agent associates said context information with said request information, respectively.

According to yet another aspect, there is provided in a multi-tier

20 computing environment including a plurality of tiers, a system for application performance management, associated with at least two selected monitored tiers of said plurality of tiers, the system comprising: for each of at least two of said monitored tiers, a respective dynamic tier extension being coupled with at least a request entry port and a request exit port of said monitored tier, said dynamic tier

25 extension at least monitoring request traffic passing through said respective monitored tier, said monitored request traffic including at least one entering request received at said request entry port and at least one respective exiting request at said request exit port, said selected tier providing said at least one respective exiting request to an adjacent respective monitored tier, said dynamic tier extension

30 capturing a request context of said at least one entering request, said dynamic tier

77964-10

extension assigning a unique transaction identifier to said at least one entering request for uniquely identifying each request of a transaction to which said at least one entering request belongs; and for each of said at least two of said monitored tiers, a respective context agent, coupled with said respective dynamic tier extension, said context agent being coupled with other said context agents associated with tiers that are directly coupled with said respective tier, said respective context agent receiving request information from said respective dynamic tier extension, said request information at least identifying each request included in said monitored request traffic, said respective context agent further receives context information relating to the request context of said at least one entering request from said respective dynamic tier extension, for each said at least one entering request, said respective context agent associates said at least one entering request with said at least one exiting request, and said respective context agent provides context information relating to the request context of said at least one respective exiting request directly to another context agent associated with said adjacent respective monitored tier, said context information including at least said unique transaction identifier.

According to a further aspect, there is provided in a multi-tier computing environment including a plurality of tiers, an application performance management method for monitoring a selected monitored tier of said plurality of tiers, the method comprising the procedures of: receiving at a context agent associated with said selected monitored tier context information relating to a request context of at least one entering request from another context agent associated with another tier which sent said at least one entering request to said selected monitored tier, said context information including at least a request identifier and a unique transaction identifier for uniquely identifying each request of a transaction to which said at least one entering request belongs; monitoring request traffic passing through said monitored tier, said monitored request traffic including at least said at least one entering request; identifying said at least one entering request in accordance with said request identifier; and associating said at least one entering request with a transaction to

77964-10

which said at least one entering request belongs in accordance with said transaction identifier.

According to yet a further aspect, there is provided in a multi-tier computing environment including a plurality of tiers, an application performance management method for monitoring a selected monitored tier of said plurality of tiers, the method comprising the procedures of: monitoring request traffic passing through said monitored tier, said monitored request traffic including at least one entering request and at least one respective exiting request, said selected tier providing said at least one respective exiting request to an adjacent tier; determining context information relating to the request context of said at least one entering request, said context information including at least a unique transaction identifier for uniquely identifying each request of a transaction to which said at least one entering request belongs; identifying each request included in said monitored request traffic; associating said at least one entering request with said at least one exiting request; and sending context information relating to the request context of said at least one exiting request directly to another context agent associated with said adjacent tier.

BRIEF DESCRIPTION OF THE DRAWINGS

The disclosed technique will be understood and appreciated more fully from the following detailed description taken in conjunction with
5 the drawings in which:

Figure 1 is a schematic illustration of a multi-tier computing environment, which is known in the art;

Figure 2 is a schematic illustration of an application performance management system, constructed and operative in accordance with an
10 embodiment of the disclosed technique;

Figure 3 is a schematic illustration of information relaying between two of the context agents of the system of Figure 2;

Figure 4 is a schematic illustration of a sample request life cycle over two of the tiers of the system of Figure 2;

15 Figure 5 is a schematic illustration of an application performance management system, constructed and operative in accordance with another embodiment of the disclosed technique;

Figure 6 is a schematic illustration of two of the tiers of the multi-tier computing environment of Figure 5;

20 Figure 7 is a schematic illustration of a dynamic tier extension of the system of Figure 5;

Figure 8 is a schematic illustration of a context agent of the system of Figure 5;

002412CA

Figure 9 is a block diagram demonstrating the stages involved in capturing a request context and subsequent processing, operative in accordance with a further embodiment of the disclosed technique;

Figure 10 is a block diagram demonstrating the stages involved
5 in capturing a UOW allocation on a local tier of the system of Figure 5 and associating a request with the UOW, operative in accordance with yet another embodiment of the disclosed technique; and

Figure 11 is a block diagram demonstrating the stages involved
in capturing an exiting request sent to a remote tier of the system of Figure
10 5, and associating the sent request with the request context, operative in accordance with yet a further embodiment of the disclosed technique.

DETAILED DESCRIPTION OF THE EMBODIMENTS

The disclosed technique overcomes the disadvantages of the prior art by providing a system and method for application performance management in a multi-tier computing environment. The system monitors
5 the request entry ports and request exit ports of each tier, and detects the entry or exit of requests to or from a given tier, via a plurality of context agents. Each context agent is associated with a tier in the multi-tier computing environment, and capable of communicating with other context agents. A context agent collects information relating to the execution of
10 requests on the tier associated therewith. The context agent identifies the request context of a user request. The context agent classifies the user request into a request class. The context agent relays characteristics of a request exiting the tier associated therewith, to a subsequent context agent associated with the tier to which the request is sent.

15 The context agent associates a request with a user request and with other previous requests in the same transaction. The context agent assigns a service class to the request in accordance with the request class and a locally stored active service class policy. The context agent may perform intervention to influence the processing of the request, such as
20 adjusting the order of the request on the queue at a request entry port to the tier, altering the priority of a processing enclave executing the request, altering the type of processing of a processing enclave executing the request, instructing the tier to allocate, or to deny, computational

resources (e.g. central processing unit - CPU, memory, and the like) to process the request, putting the request on hold and freeing the processing enclave, or terminating the request. A context network management server may profile the behavior of different types of requests across several tiers and may set an appropriate cross-tier service class policy. The system thereby provides for context related resource management at a transaction level, across the different tiers in the multi-tier computing environment.

The disclosed technique provides proactive transaction workload management capability across all tiers in an IT infrastructure chain. The system integrates with the IT infrastructure tiers, such as web, application, database, and middleware servers. The system automatically profiles workloads, helps classify the workloads, and enables a user to create appropriate service class performance policies. The system continuously applies these policies to transactions across the tiers in the computing environment. The system utilizes the existing IT infrastructure and enhances the existing IT infrastructure to enable the balanced delivery of services at optimal service levels consistent with business interests. The following terms are used throughout the description of the embodiments:

- The term “tier” herein below, refers to an entity that delivers a certain type of service, wherein the service is part of the overall delivery of an IT transaction. The service may be presentation level processing on a web server tier, application functionality on

an application server tier, data processing on a database tier, and the like. Each tier typically runs on a different host machine, although there may be more than one tier operating on a single host machine, and a single tier may include multiple components residing on more than one host machine. The host machine on which at least one tier runs, is herein below referred to as a "tier host". Examples of a tier include but are not limited to: a Java 2 Platform, Enterprise Edition (J2EE) application server instance; a cluster of J2EE application server instances; a database server instance including the access components to the database server such as Java Database Connectivity / Open Database Connectivity (JDBC/ODBC) drivers; a cluster database, and the like.

- The term "transaction" represents a single process initiated by a user, such as a stage of a business process within a business application. An example of a transaction is the placing of a bid in an online auction service or the opening of a new customer account at a financial institution. A transaction is made up of a chain of requests between tiers, starting with a user request. Therefore each request is uniquely associated with a user request (i.e., the user request of the transaction). Each transaction is identified via a unique identifier, known as a "transaction ID". It is noted that a "set of related transactions"

refers to several transactions which are interrelated (e.g., each transaction represents different stages of a single business process). The handling of a request within a transaction may take into account not only the transaction, but also the set of related transactions to which the request belongs.

- The term "request" herein below, refers to a system request from one tier to another tier, in order to provide a certain service that is part of the transaction. Each request is identified via a unique identifier, known as a "request ID". Each request results in a unit of work (UOW) on the invoked tier. Examples of a request include but are not limited to: a client web browser issuing a HyperText Transport Protocol (HTTP) request to a web server; a JAVA program issuing a Remote Method Invocation (RMI) call to an application server; a J2EE application server session bean invoking an entity bean on a remote application server (via RMI), and the like.

- The term "user request" herein below, refers to the initial request initiated by either a user or an application, which originates at a tier not monitored by the disclosed technique. The user request is the first request in the chain of requests that makes up a transaction. The chain of requests can be represented as a tree structure with the user request at the root node of the tree.

- The term “UOW” herein below refers to the application code executing in the processing enclave allocated to the applicable request on that tier (i.e., a UOW invocation). A UOW is associated with a source and a destination, may have parameters (which are directives for the application code behavior), and uses tier level resources within a single tier.
- The term "processing enclave" herein below, refers to any thread, sub-process, database session, and the like, which executes a UOW in a given tier. A request is queued in the tier until an available processing enclave is allocated and the application code (i.e., a UOW) is assigned to the processing enclave. Processing enclaves are generic execution units which in turn execute different application codes.
- The term “request context” herein below refers to a set of characteristics that are initially captured from the user request, sent to subsequent requests along the chain of requests of the transaction, and may be modified at any tier along the way. The request context enables the disclosed technique to identify, track and prioritize the resulting chain of requests as part of the single transaction initiated by a user request. The request context may include for example, the characteristics of the user who submitted the request, the characteristics of the item that is the subject of the request, the geographical location from which the

request originated, the time and date at which the request is made, the set of related transactions to which the request belongs, and the like. Certain parts of the request context may be modified at subsequent tiers. For example, the service class of the user request that is added to the request context at the first tier, may be overridden by a subsequent tier (i.e., according to another embodiment of the disclosed technique).

- The term "request class" herein below, refers to a category of transactions which share one or more pre-defined request context characteristics. For example, a "stock portfolio summary query" may be classified as a "stock portfolio summary query" request class, or may be part of a larger "stock portfolio query" request class together with another transaction, such as a "stock portfolio history query". Each request class is treated in accordance with an active service class policy. Once a request class is assigned to the user request, that request class is automatically assigned to each subsequent request in the transaction initiated by that user request.
- The term "service class" herein below, refers to a set of rankings for various parameters that indicate the level of importance for processing the request. The parameters may include: the priority to be assigned to the request, the CPU percentage to be allocated to the request, the memory to be allocated to the

request, the priority in allocating and accessing input/output (I/O) devices to the request, and the like. The service class is assigned to a request executing on a given tier by the respective context agent, in accordance with the appropriate active service class policy.

- The term "service class policy" herein below, refers to a rule which assigns a service class to a request within a request class, with respect to the tier on which the request is being processed. Each context agent contains a set of tier specific service class policies, each of which maps a service class to a request class for the specific tier associated with that context agent. A "cross-tier service class policy database" describes the set of mappings of service classes to request classes for all the tiers in the multi-tier computing environment. It is noted that a user may define a set of service class policies. Such policies are referred to herein below as "user-defined service class policies".

- The term "active service class policy" contains the request class to service class mapping that is currently in effect. Multiple service class policies are supported and a different service class policy may be scheduled at different times of system operation to reflect changing workloads or various system events, or simply as an ad hoc decision.

Reference is now made to Figure 2, which is a schematic illustration of an application performance management system, generally referenced 100, constructed and operative in accordance with an embodiment of the disclosed technique. System 100 operates on a multi-tier computing environment, generally referenced 132. Computing environment 132 includes a first client 112 running a first application, a second client 114 running a second application, a first tier 102, a second tier 104, a third tier 106, a fourth tier 108, and a fifth tier 110. The first tier 102 is a web server. The second tier 104 is an application server, application server A. The third tier 106 is another application server, application server B. The fourth tier 108 is a further application server, application server C. The fifth tier 110 is a database.

First tier 102 is coupled with first client 112, with second client 114, and with second tier 104. Second tier 104 is further coupled with third tier 106. Third tier 106 is further coupled with fourth tier 108 and with fifth tier 110. Fourth tier 108 is further coupled with fifth tier 110. The first application running on first client 112 initiates a user request R1. The second application running on second client 114 initiates a user request 118.

System 100 includes a plurality of context agents 122, 124, 126, 128 and 130, and a context network management server (CNMS) 120. In the example set forth in Figure 2, there is a single context agent associated with each tier. In particular, context agents 122, 124, 126, 128

002412CA

and 130 are associated with first tier 102, second tier 104, third tier 106, fourth tier 108 and fifth tier 110, respectively. Context agents 122, 124, 126, 128 and 130 are coupled with CNMS 120. Each context agent is also coupled with other context agents in accordance with the coupling of the tiers in computing environment 132. In particular, context agent 122 is coupled with context agent 124, context agent 124 is further coupled with context agent 126, context agent 126 is further coupled with context agent 128 and with context agent 130, and context agent 128 is further coupled with context agent 130.

First client 112 requires a service from first tier 102 and first client 112 sends a user request R1 to first tier 102. User request R1 waits on a queue at a request entry port of first tier 102. First tier 102 assigns an available processing enclave X1 to process user request R1. While processing user request R1, the application logic executing in processing enclave X1 determines that processing enclave X1 cannot complete processing user request R1 without additional information or operations to be provided by second tier 104. Therefore, first tier 102 sends a new request R2 to second tier 104, requesting the additional information or operations. Second tier 104 assigns an available processing enclave X2 to process request R2. The application logic executing in processing enclave X2 determines that processing enclave X2 requires further information or operations to be provided by third tier 106. Therefore, second tier 104 sends a new request R3 to third tier 106. Third tier 106

assigns an available processing enclave X3 to process request R3. It is noted that each of requests R1, R2, and R3 is part of a single transaction which originates from the application running on first client 112.

Processing enclave X3 completes processing. Third tier 106
5 returns a reply R3' to second tier 104, in response to the earlier request R3 from second tier 104. The application logic executing in processing enclave X2 receives the reply R3' and resumes execution. Once processing enclave X2 has completed processing, second tier 104 returns a reply R2' to first tier 102 in response to the earlier request R2 from first
10 tier 102. The application logic executing in processing enclave X1 receives the reply R2' and resumes execution. Once processing enclave X1 has completed processing, first tier 102 returns a reply R1' to user request R1, which has now completed.

Each context agent monitors the tier associated therewith at the
15 request entry ports and request exit ports of the tier (represented as small circles in Figures 2 and 3). The context agent monitors request traffic passing through the associated tier, by detecting that a request has entered or exited the associated tier. If the request is a user request (i.e., the initial request in a chain of requests), the context agent of the first tier
20 identifies the request context of the user request, classifies the user request into a request class, and assigns a service class to the user request based on the contents of the active service class policy. Each context agent has a policy cache (not shown) which contains the set of tier

specific service class policies for the tier associated with the respective context agent. CNMS 120 periodically updates each context agent with the tier specific active service class policies. If the request is not a user request, the context agent receives information relating to the request context of the request, along with additional information relating to the request (i.e., "context information"), from the context agent associated with the tier where that request originated. It is noted that the minimal information included in the context information that a context agent relays to another context agent is at least: the request ID, the transaction ID, the request class, and context related data associated with the request. Context related data may include the request context itself, or an indication (e.g., a pointer) to the request context residing in another location.

The context agent associates the received context information with the request executing on the tier. The context agent may influence the processing of the request, by the respective tier, in accordance with the service class assigned to the request. For example, the context agent may adjust the order of the request on the queue at a request entry port to the tier, or may instruct the tier to allocate, or alternatively to deny, computational resources from the tier to execute the request. If the context agent detects that a request has exited the associated tier, the context agent relays context information to another context agent associated with the tier to which the request has been sent. This other context agent associates the received context information with the relevant

request, and with the processing enclave executing the request on the tier associated with this other context agent.

It is noted that the context agent monitors the request entry ports and request exit ports of the tier, rather than extensively monitoring the activity that occurs within the tier itself (e.g., the processing enclave executing a request). As a result, system 100 does not interfere with the actual operation of a given tier or the user application code executing in the tier from a software perspective, and system 100 adds minimal additional load to the tiers.

The context agent is also coupled with the associated tier via a dynamic tier extension (DTE - not shown in Figure 2). The DTE enables the context agent to collect data relating to the execution of UOWs on that tier. The context agents may send raw data to CNMS 120 for archiving purposes. The context agents may further send to CNMS 120 statistical data for aggregated analysis. The context agents may receive information from CNMS 120 such as activity profiles (defined herein below with reference to Figure 6) and new active service class policies for the handling of different types of request classes. The context agent is elaborated upon in detail in Figure 8 herein below.

In particular, context agent 122 monitors first tier 102 and detects that user request R1 has entered first tier 102. Context agent 122 identifies the request context of user request R1 and associates user request R1 with processing enclave X1 processing the request. Context

002412CA

agent 122 classifies user request R1 into an appropriate request class. Context agent 122 determines the service class of user request R1 in first tier 102, by retrieving the appropriate active service class policy in the set of service class policies context agent 122 has stored, and assigns user request R1 the determined service class. Context agent 122 adds the assigned service class to the request context. When new request R2 exits first tier 102 toward second tier 104, context agent 122 detects that request R2 is related to user request R1. Context agent 122 then sends context agent 124 information relating to the request context of user request R1, together with the request ID, the request class, and the transaction ID associated with request R2.

Reference is now made to Figure 3, which is a schematic illustration of information relaying between two of the context agents of the system of Figure 2. Context agent 122 sends context agent 124 a message 134. Message 134 includes the request ID of request R2, the transaction ID of request R2, the request class that context agent 122 classified request R2 into, and the request context of request R2. Context agent 124 receives message 134 and determines the service class of request R2 which is to be executed on second tier 104, by retrieving the appropriate active service class policy in the set of service class policies which context agent 124 has stored. Context agent 124 assigns request R2 the determined service class. For example, the request class of request R2 is group "15". Context agent 124 retrieves the active service

002412CA

class policy that maps a service class to requests of request class "15" that are executing on second tier 104. The appropriate service class policy assigns a priority of "5" to such requests, a CPU allocation of "90", a memory allocation of "48", and I/O device access priority of "2". Context agent 124 may then influence the processing of request R2 in accordance with the assigned service class.

System 100 performs application performance management on a request context basis. System 100 identifies the requests, and the characteristics relating to each request are available at the context agent associated with the tier. These characteristics may include where that request initiated, which requests preceded the request in the transaction, and what type of request it is. For example, context agent 124 identifies that request R2 operating on second tier 104 is associated with user request R1 that was processed by first tier 102 and initiated in first client 112. Since the context agent of a given tier is aware of the request context and the request class of each request which is being executed at the respective tier, the context agent can determine the appropriate tier specific service class of the respective request based on the service class policy. CNMS 120 can set the overall management policy across several tiers, respective of different request classes, and update the context agents accordingly.

Reference is now made to Figure 4, which is a schematic illustration of a sample request life cycle, generally referenced 140, over

two of the tiers of the system of Figure 2. Sample life cycle 140 depicts the stages a request undergoes as the request is being serviced in multi-tier computing environment 132 (Figure 2). It is recalled that a request causes an invocation of a UOW, which may further generate additional requests, either internally within the same tier (by sending a request to the same tier on which the UOW is being executed), or externally (by sending requests to other tiers). Hence, a user request typically spawns a series of invocations of UOWs, each of which may be performed on a different tier. The invocations of UOWs may be synchronous (i.e., the processing enclave executing the invoking UOW waits for a reply from the invoked UOW before resuming processing) or asynchronous (i.e., the processing enclave executing the invoking UOW continues to process the invoking UOW without waiting for a reply from the invoked UOW). In both cases, the UOW on the invoked tier N+1 is dedicated to the service requested by the invoking tier N. In synchronous processing there are cases where the invoking UOW on tier N waits for the invoked UOW on tier N+1 to end (i.e., the UOW on tier N+1 is de-allocated). In other cases, the invoked UOW on tier N+1 may be referenced multiple times by the invoking tier, until the invoked UOW on tier N+1 ends.

In stage 142, a first request is sent to tier N (i.e., any representative tier) in computing environment 132. The first request results in a UOW invocation on tier N to provide a service, either for a

previous tier or for a user application. The first request waits on a queue 158 on tier N.

In stage 144, the first request exits queue 158 and is allocated a UOW, UOW-A, on tier N. A UOW allocation involves assigning an available processing enclave from one of processing enclaves 162 in tier N and dispatching the request application code to run on that processing enclave. The UOW allocation occurs once target tier resources are available and it is possible to assign the application code to an available processing enclave on tier N. In sample life cycle 140, UOW-A commences execution on tier N.

In stage 146, UOW-A issues a second request to tier N+1. Tier N+1 then invokes UOW-B to execute this request of tier N. In stage 148, tier N+1 invokes UOW-B to execute the second request sent by tier N. The second request waits on a queue 160 on tier N+1. In stage 150, the second request exits queue 160 and UOW-B is allocated to the second request. UOW-B allocation results in the assignment of an available processing enclave from one of processing enclaves 164 on tier N+1 to the application code of the UOW invocation and dispatching the request application code to run on that processing enclave. UOW-B then commences execution on tier N+1. It is noted that the invocation of UOW-B is synchronous, and thus the processing enclave processing UOW-A does not continue processing while waiting for a reply from UOW-B.

In the case that the invocation of UOW-B is asynchronous, the processing enclave processing UOW-A receives an acknowledgement from tier N+1 that the second request sent from UOW-A to tier N+1 was accepted. Upon receiving the acknowledgement, the processing enclave
5 processing UOW-A resumes execution until the processing enclave eventually returns a reply to the first request. After tier N+1 accepts the asynchronous second request, the second request waits on queue 160. The second request is subsequently read by one or more processing enclaves handling the second request, until one of these processing
10 enclaves also removes the second request from queue 160. Each processing enclave handling the second request may also turn the second request into a new user request, which may itself start another chain of requests, thereby starting a new transaction.

For example, a transaction involving an asynchronous request
15 may be a user confirming the purchase of a book in an e-commerce website. The purchase request returns the user to a screen informing that the order is being processed and the user will be notified (e.g., via email or text message). The same purchase request is simultaneously placed on a message queue where the purchase request is later processed by: a
20 processing enclave sending a final approval request to the credit card company; a processing enclave sending a purchase order request to the warehouse; a processing enclave sending an accounting request to the billing system; and the like.

In stage 152, UOW-B returns a reply to UOW-A in response to the earlier invocation of UOW-A, and execution of UOW-B is now completed. The second request has now terminated. In stage 154, UOW-A receives the reply and resumes execution, at which point UOW-B
5 is actually released by tier N+1. UOW-A may then continue executing. The duration of time between when an invoking UOW makes a request and when the invoking UOW receives a reply from the invoked UOW, is known as the "latency" period, or the response time for a given UOW request.

10 In stage 156, UOW-A completes execution and the first request terminates. It is noted that before completion, UOW-A may require the services of another tier and may invoke a further request in order to provide that service.

It is noted that after UOW-B is allocated in stage 150, and
15 commences execution on tier N+1, an unrecoverable error may occur (e.g., a program exception). In accordance with an embodiment of the disclosed technique, the context agent associated with tier N+1 will record the error and will associate the error to the transaction which commenced with the user request that invoked UOW-A on tier N, providing information
20 as to the nature of the error that occurred on tier N+1.

Reference is now made to Figure 5, which is a schematic illustration of an application performance management system, generally referenced 200, constructed and operative in accordance with another

002412CA

embodiment of the disclosed technique. System 200 operates on a multi-tier computing environment, generally referenced 248. Computing environment 248 includes a client 214, a first tier 202, a second tier 204, a third tier 206, a fourth tier 208, and a fifth tier 210. First tier 202 is a web server. Second tier 204 is an application server A. Third tier 206 is an application server B. Second tier 204 and third tier 206 both reside on a single host machine 212. Fourth tier 208 is another application server, application server C. Fifth tier 210 is a database. First tier 202 is coupled with client 214, and with host 212. Host 212 is further coupled with fourth tier 208 and with fifth tier 210. Fourth tier 208 is further coupled with fifth tier 210.

System 200 includes a plurality of dynamic tier extensions 222, 224, 226, 228 and 230, a plurality of context agents 232, 234, 236 and 238, a plurality of local logs 240, 242, 244 and 246, a context network management server (CNMS) 216, a service level objective (SLO) policy database 218, and a supervisor workstation 220. Each tier contains a dynamic tier extension (DTE). There is a context agent associated with each tier. There is a local log associated with each context agent. A context agent of a given tier is coupled with the DTE (or several DTEs) within the tier, with the local log associated with the context agent, and with other context agents in accordance with the coupling of the tiers in the computing environment. Each context agent is also coupled with CNMS 216. CNMS 216 is coupled with SLO policy database 218. Supervisor

workstation 220 is coupled with CNMS 216 and with SLO policy database 218.

In particular, first tier 202 includes DTE 222. Context agent 232 is associated with first tier 202. Local log 240 is associated with context agent 232. Second tier 204 includes DTE 224. Third tier 206 includes DTE 226. Since second tier 204 and third tier 206 both reside on host 212, there is only a single context agent 234 associated with both second tier 204 and third tier 206. It is noted that context agent 234 is directly coupled with each of DTE 224 and DTE 226. Local log 242 is associated with context agent 234. Fourth tier 208 includes DTE 228. Context agent 236 is associated with fourth tier 208. Local log 244 is associated with context agent 236. Finally, fifth tier 210 includes DTE 230. Context agent 238 is associated with fifth tier 210. Local log 246 is associated with context agent 238. Context agent 232 is coupled with context agent 234. Context agent 234 is further coupled with context agent 236 and with context agent 238. Context agent 236 is further coupled with context agent 238.

A dynamic tier extension is coupled with the tier at specific predetermined points. These predetermined points are: the request entry ports of the tier, the request exit ports of the tier, and possibly additional areas within the tier (e.g., a tier control port). A request port according to the disclosed technique is a module within a tier which manages requests, either before they are processed by the tier, or after they are processed by

the tier. Such a request port can be an interface point (i.e., entrance, exit or any other access mechanism) to a request queue at the entrance of a tier. Since a request requires service from an application code running on the tier by a processing enclave, the respective request port resides at an application level and not at a networking level. It is noted that request ports according to the disclosed technique, are not at a network level (e.g., not TCP/IP or UDP ports).

The DTE is located on the same tier host as the tier associated therewith (i.e., the DTE is located on at least one of the host machines on which the tier is running). Among the responsibilities of the DTE is capturing a request context. The DTE further monitors the request entry ports and request exit ports of a tier, to detect incoming and outgoing requests. The DTE assigns a transaction ID to a user request, and obtains the request ID of each request entering or exiting the tier. The DTE is elaborated upon in detail in Figure 7 described herein below.

The context agent maintains associations between a given request, the invoked UOW of the request, and the request context of the user request in the same transaction as the request. The context agent relays the request context assigned to each request (i.e., context related data) to other context agents handling other tiers. The context agent may relay the entire request context, or a portion of the request context. Furthermore, the context agent may relay the request context itself, or an indication (e.g., a pointer) to the request context residing in another

location. It is noted that the context agent need not necessarily reside on the same host machine as the tier, but this is the case in a preferred embodiment of the disclosed technique. The context agent is elaborated upon in detail in Figure 8 described herein below. CNMS 216 collects and
5 analyzes performance data. SLO policy database 218 stores cross-tier service class policies, and is continuously updated.

Reference is now made to Figure 6, which is a schematic illustration of two of the tiers of the multi-tier computing environment of Figure 5. It is noted that each of fourth tier 208 and fifth tier 210 illustrated
10 in Figure 6 is representative of any two consecutive tiers (e.g., tier N and tier N+1) in computing environment 248.

A request entering fourth tier 208 waits at a request entry port of fourth tier 208 at a queue 262. Fourth tier 208 invokes a UOW to execute the request. The request exits queue 262 and fourth tier 208 allocates the
15 UOW to the request, by assigning an available processing enclave to the UOW from processing enclaves 252 and dispatching the request application code to run on the processing enclave. The UOW executes on fourth tier 208. The UOW may then request a service from fifth tier 210. The new request exits fourth tier 208 at a request exit port and waits at a
20 request entry port of fifth tier 210 at a queue 264. Fifth tier 210 invokes a UOW to execute the new request. The new request exits queue 264 and fifth tier 210 allocates the UOW to the new request, by assigning an available processing enclave to the UOW from processing enclaves 254,

and dispatching the application code of the new request to run on the processing enclave. DTE 228 monitors the request entry ports and request exit ports of fourth tier 208 to detect the requests entering and exiting fourth tier 208, respectively.

5 The dynamic tier extensions are involved in tracking a request throughout the life cycle thereof, without changing the application code. The DTE dynamically hooks into the tier environment where the DTE intercepts the request context external to the application code. Request tracking includes capturing the request context, associating the request to
10 a UOW on a tier, and disassociating the request from a UOW on a tier. The DTE further collects performance, availability, and error metrics from the tier. The DTE also may dynamically adjust the processing of requests on the tier, such as by adjusting the order of a request on the queue at a request entry port of the tier, allocating computational resources to
15 process the request (e.g., CPU, memory, I/O, and the like) or altering the priority of the processing enclave or the allocated resources. These tasks are elaborated upon with reference to Figures 9, 10 and 11 described herein below.

It is noted that there are two alternatives for request context
20 relaying between context agents: in-band and out-of-band. In in-band context relay, the request context is added to the request itself (i.e., onto the payload), as the request exits a certain tier toward the next tier. Therefore, as a request and descendant requests thereof are being

processed between different tiers, the updated request context is added to the request invocations. In contrast, out-of-band relaying does not involve the request context being added to the request invocations. Rather the context agents send the context information directly to each other. A context agent sends a request context to another context agent. A DTE retrieves the request context from the context agent. It is noted that each context agent of a system similar to system 200 (Figure 5) relays the request context to another context agent using the out-of-band technique.

With reference back to Figure 5, system 200 performs activity profiling. Activity profiling involves creating an activity profile. The activity profile includes integrated transaction, tier, and system level performance metrics and statistical analysis, which are obtained during request tracking. For example, the performance metrics may include: request elapsed time, request service time, consumed CPU time on the tier, and the like. The activity profile data is collected over time and used for both monitoring (i.e., display in the GUI) and for supporting user creation of a user-defined service class policy. An activity profiling console (not shown) is a component of the GUI that displays performance and availability aggregated data gathered by the context agent. The performance and availability aggregated data includes summary views of activity profile data by various categories such as: tier, request class, transaction and the like. An activity profiling engine (not shown) located in CNMS 216 performs activity profiling.

Each context agent has the task of collecting information about the execution details in each tier. The DTE enables the context agent to collect data relating to the execution of UOWs on that tier. The context agents then send the collected information to CNMS 216. The information
5 stored for each UOW includes: start time, request ID of the request to which the UOW is allocated, transaction ID of the request to which the UOW is allocated, request class of the request to which the UOW is allocated, user details, originating network address details, service class of the request to which the UOW is allocated, end time, resource
10 consumption (such as a CPU), and the like.

The context agent stores the information relating to currently executing UOWs in a memory (not shown). Once the UOW has finished executing on the tier the context agent transfers the information to a recent history data store (not shown), which is stored on a disk (not shown)
15 locally at the same tier host of the context agent. After a certain period, the context agent moves the entries from the recent history data store to a summary data store (not shown). This information is stored as a summary over a given period of time (e.g., an average of collected metrics over a period of half an hour). The information in the summary data store is
20 stored in shifts. Shifts are saved on a periodic basis (i.e., they are recycled after a number of shifts have accumulated).

System 200 further includes service class policies. It is recalled that a service class policy refers to a rule which assigns a service class to

a request within a request class, with respect to the tier on which the request is being processed. A real time service class policy engine, located in the context agent, assigns an appropriate service class to the request in accordance with information in the request context, such as the request class, performance characteristics of previous executions of requests with the same request context, and in accordance with an active service class policy. Furthermore, assignment of a service class may take into account additional considerations, such as a set of performance metrics for each request class. This set of performance metrics characterizes the request class and creates a baseline for typical performance behavior. This process utilizes the log data created across all tiers and the request classes associated with them. The request classes are stored in a tree structure in SLO policy database 218. The service class policy engine adds and updates information relating to each request class, and updates SLO policy database 218 accordingly.

The service class is determined for the user request in the first tier and then passed from tier to tier along with the request in the request context (from context agent to context agent). There is no need to access CNMS 216 in order to determine the active service class policy and the mapping to the service class. Rather, each context agent has a policy cache therein (not shown) kept up to date, so that the mapping to the appropriate service class is performed locally.

The activity profiling engine polls the local context agents for information about completed and currently running requests at periodic intervals, or whenever necessary and on demand (e.g., when a user wishes to view the currently running requests on the performance console). Alternatively, the activity profiling engine can instruct each context agent to initiate sending new records to the activity profiling engine at fixed intervals or when a threshold (i.e., number of records) is reached. The activity profiling engine collects the data for each request, performs calculations (e.g., average, variance, and the like) and stores the results at various intervals in SLO policy database 218, as a baseline for analysis for each request class. Based on the available data stored in SLO policy database 218, a service class policy generation engine located in CNMS 216 creates a set of rules that serve as recommendations for new service class policies. CNMS 216 determines service class policies using these recommendations. The service class policies are stored in SLO policy database 218.

It is noted that SLO policy database 218 stores policies that are automatically generated as well as user-defined service class policies. A user may create a set of service class policies (i.e., user-defined service class policies) via the user interface (not shown) of system 200, or by editing a configuration file of the context agent. The creation of user-defined service class policies involves user analysis of the activity

profiles, and obtaining approval of the suggested service class policy by CNMS 216.

The context agent receives updated service class policies from CNMS 216. It is noted that whether the new service class policy is automatically generated or user-defined is transparent to the context agent. The context agent assigns to a request the service class designated in the appropriate tier-specific active service class policy, located within the local policy cache of the context agent.

Alternatively, the context agent may assign the request a different service class than that designated by the active service class policy in certain situations (e.g., if the tier is currently unable to provide all the resources required to fulfill the service class, if a significant amount of high priority requests enter the tier and may result in resource starvation for lower priority requests, and the like). Further alternatively, the context agent may alter the request class of the request, and subsequently assign the request the service class designated by the appropriate tier-specific active service class policy for the new request class.

System 200 also performs request classification. This involves classifying multiple requests (each designated by their transaction ID and request context) into request classes according to various criteria. An automatic request classification process gathers the generic information associated with a request (e.g., header information, parameters such as query string, Uniform Resource Locator (URL) parameters in the case of

an HTTP request type, and the like). When another request of the same request class arrives, this request will be treated in a similar manner as other requests within the request class. The request classes may be grouped into various categories in accordance with the request class performance characteristics in the relevant activity profiles (e.g., a group of request classes having a response time greater than two seconds). Request class grouping is used for reporting purposes, enabling higher level summarized views of request class performance data.

System 200 also supports a user defined classification process wherein the user creates rules that classify requests based on the same request information used for the automatic classification process. A user may create classification rules via the user interface (not shown) of system 200. The automatic request classification process uses a "class based cache management" algorithm, as described by H. Zhu and T. Yang (*"Class-based cache management for dynamic web contents,"* Tech. Rep. TRCS00-13, Dept. of Computer Science, University of California, Santa Barbara, 2000). The output of the classification process is a tree representing the request class and the parent classes (i.e., in an automatic request classification process) or a set of classification rules (i.e., in a user defined classification process). All the results of classification processes are stored in the SLO policy database 218.

System 200 is also operative to enforce the service class, as defined in the active service class policy, on each and every monitored

tier. Policy enforcement is performed by both the DTE and the context agent. Enforcement may be implemented either by controlling the queuing of requests in each tier (i.e., the order in which the request is actually processed within the tier), or by temporarily changing the processing
5 priority of the processing enclave running the request during execution.

The implementation depends on the particular tier architecture. For example, implementation in a database tier may involve use of a database resource management scheme to manipulate session resource allocation according to the appropriate service class policy. Another
10 example is implementation in an application server tier implemented using a J2EE Application Server application which may involve: extending the application server web queue to support request prioritization, extending the Enterprise JavaBeans (EJB) queue to support prioritization, controlling the JAVA thread priority, and the like.

15 Reference is now made to Figure 7, which is a schematic illustration of a dynamic tier extension of the system of Figure 5. It is noted that DTE 228 of fourth tier 208, depicted in Figure 5, is representative of all the dynamic tier extensions in system 200. DTE 228 includes a series of soft hooks or interception points into fourth tier 208.
20 These hooks, referenced 266 and 268, serve to collect request related information and performance data. The hooks can also alter the priority of a processing enclave executing a UOW. A hook places the collected information on a dedicated message queue 270. The soft hooking and

interception technique depends on the particular environment. For example, the environment may be: a web server, a JAVA based J2EE application server, a database, a messaging server, and the like. It is noted that interception points 266 and 268 may be instantly activated or
5 deactivated by an operator.

DTE 228 further includes a process running within fourth tier 208. The process handles messages from the interception points, communicates messages to context agent 236, returns messages to the interception points, and performs administrative control functions of the
10 DTE (e.g., start/stop tracking requests, installing and removing soft hooks). DTE 228 includes a DTE daemon 272, a DTE administrator 274, a DTE messaging module 276, a DTE log event handler 278 and a DTE communication interface 280.

DTE daemon 272 is an artificially created processing enclave
15 operating within the tier. DTE daemon 272 performs asynchronous processing associated with reported events where the request does not need to be detained. There are two types of scenarios concerning reporting events. In the first scenario, there is no need to detain the request until a response is received. For example, when reporting that a
20 request has ended, there is no need to delay the request until after the context agent has actually been notified. In the second scenario, the request needs to be held for a certain period before processing may resume. For example, when obtaining the service class of a request or

when performing classification of a request, processing cannot commence until the request class and service class is determined by the context agent, otherwise processing may be done using an incorrect service class or an incorrect request class.

5 DTE daemon 272 deals with events of the first scenario, where the request does not need to be detained. The processing is done asynchronously, such that the request is not delayed. The request is thereby released very quickly, almost instantaneously. DTE daemon 272 has a queue 270 associated therewith. After request entries and related
10 request exits are notified by interception points 266 and 268, DTE daemon 272 picks up these notifications from queue 270 and performs whatever additional processing is necessary.

 DTE administrator 274 enables DTE 228 to receive messages relating to how requests should be processed. For example, such
15 messages may include: stop tracking a request, continue tracking the request but stop prioritizing, and the like. DTE messaging module 276 communicates with context agent 236 using messages. For example, such messages may include: start or end of a UOW, associate a UOW with a given request, and the like. DTE log event handler 278 logs tracing
20 information concerning DTE operations and logs alerts raised by DTE 228. These logged events could be sent to multiple destinations such as a local file, system message console, system log, and the like. DTE log event handler 278 supports multiple industry standard protocols such as Simple

Network Management Protocol (SNMP), and the like. DTE communication interface 280 serves as an interface between DTE 228 and context agent 236. DTE communication interface 280 relays messages sent from agent messaging module 286 of context agent 236 to DTE 228. DTE communication interface 280 also relays messages sent from DTE messaging module 276 of DTE 228 to context agent 236. Multiple communication protocols are supported, and each DTE uses the most efficient communication method available within its architecture, such as inter process communication, Transmission Control Protocol/Internet Protocol (TCP/IP), and the like.

Reference is now made to Figure 8, which is a schematic illustration of a context agent of the system of Figure 5. It is noted that context agent 236 of fourth tier 208, depicted in Figure 5, is representative of all the context agents in system 200. Context agent 236 receives notifications from DTE 228 via a variety of mechanisms, such as TCP/IP, inter-process communication channels, and the like. The notifications are of events occurring within the tier, such as: the capture of a request context, the start of an UOW, the end of an UOW, resource consumption of an UOW, the invocation/allocation of an UOW on a remote tier, the reply/release of an UOW on a remote tier, and the like. Context agent 236 includes an agent log event handler 282, an agent communication interface 284, an agent messaging module 286, a context table manager

288, a classification manager 290, a policy manager 292, and an agent administrator 294.

Agent log event handler 282 is used by context agent 236 both for internal housekeeping purposes and for logging alerts raised by context agent 236. Agent log event handler 282 logs information that enters the context table (as described herein below), but is also used for internal tracing and messaging purposes, such as to detect operational irregularities (i.e., problems or errors) that may occur within context agent 236. These logged events can be sent to multiple destinations such as a local file, system message console, system log, and the like. Agent log event handler 282 supports multiple industry standard protocols such as SNMP, and the like.

Agent communication interface 284 serves as an interface between DTE 228 and context agent 236. Agent communication interface 284 relays messages sent from DTE messaging module 276 of DTE 228 to context agent 236. Agent communication interface 284 also relays messages sent from agent messaging module 286 of context agent 236 to DTE 228. There may be several channels connecting DTE 228 and context agent 236, to ensure fast and reliable communication between the two, and at least several channels are kept open at all times, for example a high priority channel and an administrative channel. There may also be several connections of each type of channel, for different types of

messages. As a result, agent communication interface 284 is operative to accommodate these different possibilities.

Agent messaging module 286 notifies other context agents associated with remote tiers that a request was sent to the remote tier.

5 Agent messaging module 286 further communicates with DTE 228 using messages. For example, such message include: start or end of a UOW, associate a UOW with a request, and the like. Agent messaging module 286 communicates with DTE 228 via agent communication interface 284.

Context table manager 288 functions as the bookkeeper of the
10 context agent. Context table manager 288 maintains a cross-reference table, known as a "context table", used to associate UOWs running on the tier to their request context. The request context may be from the current tier (i.e., in the case of a user request), relayed from a previous tier from the request execution chain, or both (i.e., the request context is modified
15 or new information is added therein). The context table stores information associated with each request (e.g., transaction ID, request class, service class, origin of request, and the like). Agent messaging module 286 accesses the context table and looks up a record of interest after agent messaging module 286 has received information from DTE 228. Context
20 table manager 288 identifies a request based on information associated with the request and the data stored in the context table. In this manner, context agent 236 obtains information relating to the request entering the

tier, such as the request class, service class, and other relevant information associated with the request.

Classification manager 290 and policy manager 292 process each user request that enters the first tier. The first time a user request enters the computing environment there is no information in the context table concerning this user request. The user request is therefore required to undergo classification. During the classification process, all the information known about the user request at a request entry port is collected. For example, if the user request is an HTTP request, then such information includes the HTTP header, the query string, the URL parameters, and the like. For each type of protocol used in each of the tiers, there is a generic plug-in that actually classifies the user request.

The classification process essentially extracts the business perspective from the user request, by translating technical request related information (e.g., an HTTP request) into a formal request classification related to a business process (e.g., retrieve bank account balance). A user request is placed into a specific request class. The user request may be identified as part of a business process or a set of related transactions. For example, retrieval of an account balance may be part of a larger process of applying for a mortgage. When information is transferred from context agent 236 to CNMS 216, CNMS 216 can determine the activity profile and detect request behavior trends for request classes.

Policy manager 292 assigns a service class to the requests. Policy manager 292 receives the output of classification manager 290, and based on the request class, the request context, and the active service class policy, determines the service class of a given request. For
5 example, policy manager 292 may set all requests from a certain request class to have a service class with a higher priority than all requests from a different request class. A context agent assigns a service class to a request being processed on the tier associated therewith, by updating the request context with the appropriate service class according to the active
10 service class policy.

The cross-tier service class policy database includes the set of mappings of service classes to request classes for all the tiers in the multi-tier computing environment. Each request belonging to a certain request class may be assigned a different service class depending on the tier in
15 which the request is being processed. With reference to Figure 2, user request R1 may be assigned a service class with a low priority on first tier 102, request R2 may be assigned a service class with a high priority on second tier 104 and request R3 may be assigned a service class with a medium priority on third tier 106. The cross-tier service class policies are
20 stored in SLO policy database 218. Cross-tier service class policies may be either automatically generated (i.e., system defined) or defined by a user of the system.

In addition, a supervisor of a local tier (e.g., a database administrator) has administrative control of the tier and may decide to override a service class policy on that tier if deemed necessary. The tier supervisor may alter the service class assigned to a request by the context agent of that tier. The tier supervisor has an overall view of substantially everything running on the tier. It is noted that there may also be an overall supervisor of the system, who is typically a person who runs the user application itself and is interested in the global picture.

It is noted that after the initial classification and policy assignment of a user request, the request class is maintained for subsequent requests of the same transaction, as the context agents relay this information along from tier to tier (i.e., stored in the context table of each tier). In this way, at each tier the context agent identifies which request class a request belongs to, what the tier-specific service class of the request is, and other information associated with the request.

Agent administrator 294 is an interface to CNMS 216. Agent administrator 294 reports historical data to CNMS 216. For example, when a DTE indicates that a UOW has ended, the information associated with that UOW need no longer remain in the context table. The information is then sent from the context table to agent administrator 294, which archives the information and periodically sends the information to CNMS 216. Agent administrator 294 also receives from CNMS 216 new active service class policies, new configurations of the context agent, and

the like. Agent administrator 294 may also be interrogated in real-time to obtain a status indication. The status indication may include what information is currently in the context table, which UOWs are now running, and the like.

5 Reference is now made to Figure 9, which is a block diagram demonstrating the stages involved in capturing a request context and subsequent processing, operative in accordance with a further embodiment of the disclosed technique. In procedure 310, a request context is captured. A new user request is identified and the request
10 context thereof is intercepted by the DTE soft hook. With reference to Figure 5, DTE 222 captures a request context of a user request that has entered first tier 202.

 In procedure 312, the captured request context is normalized. The DTE converts the request context to a standard format, identical for all
15 types of tiers (e.g., databases, application servers, and the like). The DTE assigns the user request a unique identification (i.e., transaction ID) that will identify the user request and subsequent requests in the transaction initiated by the user request. The DTE further obtains the request ID of the user request from the tier. With reference to Figure 5, DTE 222
20 converts the captured request context to a standard format, assigns a transaction ID to the user request, and obtains the request ID of the user request.

In procedure 314, the DTE sends the request context, the transaction ID, and the request ID to the context agent associated with that tier. The DTE notifies the context agent that a new request context has been sent. With reference to Figure 5, DTE 222 sends the request
5 context, the transaction ID, and the request ID of the user request to context agent 232.

In procedure 316, the user request is classified. The context agent applies a tier specific classification scheme that determines the request class based on the request context. With reference to Figure 5,
10 context agent 232 classifies the user request into a certain request class.

In procedure 318, the appropriate active service class policy is retrieved from the service class policies stored in the local policy cache. The context agent retrieves the active service class policy using the request class, and possibly other request context fields. With reference to
15 Figure 5, context agent 232 retrieves the appropriate active service class policy for the user request, in accordance with the request class of the user request, and specific to first tier 202.

In procedure 320, the service class for the user request is determined according to the retrieved service class policy, and assigned to
20 the request. The context agent may then add the assigned service class to the request context. With reference to Figure 5, context agent 232 assigns the user request the service class of the retrieved appropriate service class policy, and adds the service class to the request context.

In procedure 322, the tier priority is determined for the UOW invoked by the user request. A local priority for the processing enclave executing the UOW is extracted from the service class. With reference to Figure 5, context agent 232 extracts the local priority for the processing
5 enclave allocated to the user request for executing the UOW, from the assigned service class.

In procedure 324, the context agent sends back the request class and prioritization information to the DTE. This information is necessary for certain tasks, such as altering the priority of a processing
10 enclave, which takes place within the DTE. With reference to Figure 5, context agent 232 sends the assigned request class and the local priority for the processing enclave executing the UOW to DTE 222.

In procedure 326, the request context and related information (such as request class and transaction ID) is forwarded to other context
15 agents and to internal components of the context agent, most notably the context table that stores the request, which indexes the request context and related information for further reference. With reference to Figure 5, context agent 232 forwards the request context, along with additional information relating to the request (e.g., the request ID, transaction ID, and
20 request class) to other context agents coupled therewith (e.g., context agent 234), as well as to internal components of context agent 232, such as context table manager 288, classification manager 290, and policy manager 292 (with reference to Figure 8).

Reference is now made to Figure 10, which is a block diagram demonstrating the stages involved in capturing a UOW allocation on a local tier of the system of Figure 5 and associating a request with the UOW, operative in accordance with yet another embodiment of the disclosed technique. It is recalled that a UOW is the application logic executing in the processing enclave associated with the request on that tier. In procedure 340, a UOW allocation is captured by the DTE. The capturing may occur when the UOW is initially invoked by the tier, as an entering request enters the queue at a request entry port of the tier, while the entering request remains queued, or once resources are available at the tier and the request exits the queue. With reference to Figure 5, DTE 228 captures a UOW allocation by fourth tier 208.

In procedure 342, the DTE determines the request ID of the entering request associated with the UOW allocation. The request ID is determined based on the information sent from the previous tier, including the request context together with an association key that links the UOW to the request, such as socket identifiers, and the like. With reference to Figure 5, DTE 228 determines the request ID of the entering request associated with the captured UOW allocation, based on information received from a previous coupled tier (e.g., third tier 206).

There might be situations where the request context and association key do not reach the DTE at the time procedure 342 takes place. Therefore the system could be configured to proceed either by

waiting for this information or by assigning a temporary request ID and associating the entering request with a UOW allocation in a later stage.

In procedure 344, the DTE sends information relating to the UOW identification as well as the determined request ID of the entering
5 request, to the context agent. A UOW identification is a set of characteristics uniquely identifying the processing enclave executing the UOW. The UOW identification is used by the context agent for keeping track of the UOW. With reference to Figure 5, DTE 228 sends association information for the UOW (e.g., an association key that links the UOW to
10 the request, such as socket identifiers), as well as the determined request ID of the entering request, to context agent 236.

In procedure 346, the context agent retrieves the entry in the context table associated with the UOW identification or the request ID. With reference to Figure 5, context agent 232 locates the entry in the
15 context table via context table manager 288 (Figure 8) associated with the captured UOW allocation or the determined request ID.

If the entry is found, then in procedure 348, that entry is updated with the UOW identification and related information. With reference to Figure 5, context agent 232 updates the relevant entry in the context table
20 via context table manager 288 (Figure 8) with the identification of the captured UOW allocation and related information.

If the entry is not found, then in procedure 350, a new entry is added to the context table. The new entry includes the request ID of the

entering request and the identification of the associated UOW allocation. The default request class and service class are associated with the newly added entry. With reference to Figure 5, context agent 232 adds a new entry to the context table via context table manager 288 (Figure 8) that
5 includes the request ID of the entering request and identification of the UOW allocation, and associates a default request class and service class to the newly added entry.

In procedure 352, the context agent determines the local priority for the processing enclave executing the UOW, and gathers request
10 statistics if necessary. With reference to Figure 5, context agent 232 extracts the local priority for the processing enclave executing the UOW invoked by the request from the service class assigned to the request.

In procedure 354, the context agent enforces the service class assigned to the entering request. With reference to Figure 5, context
15 agent 232 influences the processing of the entering request at first tier 202, by for example, altering the priority level of the processing enclave executing the UOW invoked by the request, altering the type of execution of the processing enclave, or allocating or denying computational resources to process the request.

20 It is noted that, the context agent can subsequently modify the service class, request class, or other parameters of the request context, if the context agent receives new association information relating to the request. This occurs in a situation where the request context arrives at the

context agent associated with the tier from a remote context agent (e.g., with reference to Figure 5, context agent 234 receives information about a request from context agent 232). Due to timing issues, a context agent may capture the UOW allocation associated with a request, and then
5 receive the request context of the request at a later stage.

Reference is now made to Figure 11, which is a block diagram demonstrating the stages involved in capturing an exiting request sent to a remote tier of the system of Figure 5, and associating the sent request with the request context, operative in accordance with yet a further
10 embodiment of the disclosed technique. In procedure 370, a request sent to a second tier from a first tier is captured by the DTE of the first tier. With reference to Figure 5, a request is sent from first tier 202 to second tier 204. DTE 222 captures the exiting request on first tier 202.

In procedure 372, the DTE determines the request ID of the
15 exiting request. With reference to Figure 5, DTE 222 determines the request ID of an exiting request sent from first tier 202 to second tier 204.

In procedure 374, the DTE sends information relating to the UOW identification, as well as the determined request ID of the exiting request, to the local context agent. With reference to Figure 5, DTE 228
20 sends information relating to the UOW identification, as well as the determined request ID of the exiting request, to context agent 236.

In procedure 376, the local context agent retrieves the entry in the context table associated with the request ID of the exiting request or

the UOW identification. With reference to Figure 5, context agent 232 locates the entry in the context table via context table manager 288 (Figure 8) associated with the determined exiting request ID or by the UOW identification.

5 If the entry is found, then in procedure 378, that entry is updated with the UOW identification and related information. For example, the entry may be updated with information in the request context of the exiting request that is not currently in the entry. With reference to Figure 5, context agent 232 updates the relevant entry in the context table via
10 context table manager 288 (Figure 8) with the UOW identification and with information in the request context of the exiting request that was not previously stored in the relevant entry.

 If the entry is not found, then in procedure 380, a new entry is added to the context table. The new entry includes the request ID of the
15 exiting request and the identification of the associated UOW allocation. The request class, service class, and other characteristics of the request stored in the request context are added to the new entry. If certain characteristics are not present in the request context, then default characteristics (e.g., request class and service class) are associated with
20 the newly added entry. With reference to Figure 5, context agent 232 adds a new entry to the context table via context table manager 288 (Figure 8) that includes the request ID of the exiting request and

characteristics of the request stored in the request context of the exiting request.

In procedure 382, the local context agent sends a message to the remote context agent (i.e., the context agent to which the exiting request was sent). The message includes the request ID of the exiting request, as well as the transaction ID and request class. With reference to Figure 5, context agent 232 associated with first tier 202 sends a message to context agent 235 associated with second tier 204. The message includes the request ID of the exiting request, as well as the transaction ID and the request class thereof.

In a further embodiment of the disclosed technique, there may be a single context agent associated with multiple tiers hosts. It is noted that the manner in which a context agent communicates with a DTE, enables the context agent to reside anywhere on the network. It is further noted that the context agent keeps track of request and UOWs by attaching a tier identifier thereto, thereby providing for separate bookkeeping for each tier. Such an embodiment may be used, for example, due to security concerns or concerns of extra overhead of a user resulting from the addition of the context agent to the tier host.

It will be appreciated by persons skilled in the art that the disclosed technique is not limited to what has been particularly shown and described hereinabove. Rather the scope of the disclosed technique is defined only by the claims, which follow.

77964-10

CLAIMS:

1. An apparatus for monitoring a selected tier in a multi-tier computing environment, said selected tier is coupled with at least one adjacent tier of said multi-tier computing environment, said apparatus comprising:

5 a context agent associated with said selected tier of said computing environment, said context agent being coupled with other context agents, each of said other context agents associated with a respective one of said at least one adjacent tier; and

a dynamic tier extension, said dynamic tier extension coupled with said
 10 context agent and with at least a request entry port of said selected tier, at least for monitoring request traffic passing through said selected tier, said monitored request traffic including at least one entering request received at said request entry port from a sending tier of said at least one adjacent tier, said dynamic tier extension providing request information to said context agent, said request information enabling said
 15 context agent to identify each request included in said monitored request traffic,

wherein said context agent receives context information directly from another context agent associated with said sending tier of said at least one adjacent tier, said received context information relating to the request context of said at least one entering request, said received context information including at least a unique
 20 transaction identifier assigned to each request of a transaction to which said at least one entering request belongs; and

wherein for each said at least one entering request, said context agent associates said context information with said request information, respectively.

2. The apparatus according to Claim 1, wherein said dynamic tier
 25 extension is further coupled with at least a request exit port of said selected tier, and

77964-10

wherein said monitored request traffic further includes at least one exiting request received at said request exit port, said selected tier providing said at least one exiting request to receiving tier of said at least one adjacent tier of said computing environment,

5 wherein said context agent associates between said at least one exiting request and said at least one entering request; and sends context information relating to the request context of said at least one exiting request to a further context agent associated with said receiving tier of said at least one adjacent tier.

3. The apparatus according to Claim 1, wherein said dynamic tier
10 extension is further coupled with a tier control port of said selected tier.

4. The apparatus according to Claim 3, wherein said context agent further alters the priority level of a processing enclave executing said at least one entering request.

5. The apparatus according to Claim 3, wherein said context agent further
15 alters the type of execution of a processing enclave executing said at least one entering request.

6. The apparatus according to either of Claims 1 or 3, wherein said context agent further adjusts the order of said at least one entering request on a queue at said request entry port.

20 7. The apparatus according to either of Claims 1 or 3, wherein said context agent further instructs said selected tier to allocate computational resources to process said at least one entering request.

8. The apparatus according to either of Claims 1 or 3, wherein said context agent further instructs said selected tier to deny computational resources to
25 process said at least one entering request.

77964-10

9. The apparatus according to Claim 1, wherein said context information includes at least a request identifier, a request class, and context related data.

10. The apparatus according to Claim 9, wherein said context related data includes said request context, of said at least one entering request, in its entirety.

5 11. The apparatus according to Claim 9, wherein said context related data includes a portion of said request context of said at least one entering request.

12. The apparatus according to Claim 9, wherein said context related data includes an indication to said request context of said at least one entering request.

13. The apparatus according to Claim 2, wherein said context agent
10 comprises:

an agent communication interface, relaying messages between said dynamic tier extension and said context agent;

an agent messaging module, notifying said further context agent that said at least one exiting request has been sent to said another adjacent tier, said
15 agent messaging module further communicating with said dynamic tier extension using messages; and

a context table manager, maintaining a reference table associating at least one unit of work running on said selected tier with said request context of said at least one entering request to which said unit of work has been allocated, said
20 reference table being a context table, said context table further storing information associated with said at least one entering request, said context table manager further identifying at least one entering request based on information stored in said context table.

14. The apparatus according to Claim 13, wherein said context agent
25 further comprises:

77964-10

a classification manager, performing a classification of said at least one entering request based on said received context information, said classification manager further identifying said at least one entering request as part of said transaction.

- 5 15. The apparatus according to Claim 14, wherein said context agent further comprises:

a policy manager, assigning a service class policy to said at least one entering request according to an active service class policy and according to said classification of said at least one entering request.

- 10 16. The apparatus according to Claim 13, wherein said context agent further comprises:

an agent log event handler, performing event logging, internal tracing and messaging, and detection of operational irregularities within said context agent.

- 15 17. The apparatus according to Claim 13, wherein said context agent further comprises:

an agent administrator, archiving information, sending historical data to a context network management server (CNMS) and receiving active service class policies and additional data from said CNMS.

- 20 18. The apparatus according to Claim 3, wherein said dynamic tier extension comprises:

a dynamic tier extension daemon, performing asynchronous processing associated with reported events where said at least one entering request is not detained;

- 25 a dynamic tier extension messaging module, communicating with said context agent using messages; and

77964-10

a dynamic tier extension communication interface, relaying messages between said dynamic tier extension and said context agent.

19. The apparatus according to Claim 18, wherein said dynamic tier extension further comprises:

5 a dynamic tier extension administrator, enabling said dynamic tier extension to receive messages relating to how said at least one entering request should be processed.

20. The apparatus according to Claim 18, wherein said dynamic tier extension further comprises:

10 a dynamic tier extension log event handler, logging tracing information concerning operation of said dynamic tier extension and logging alerts raised by said dynamic tier extension.

21. The apparatus according to Claim 1, wherein said context agent resides on the same host machine as said selected tier.

15 22. The apparatus according to Claim 3, wherein said dynamic tier extension collects performance, availability, and error metrics from said selected tier.

23. An apparatus for monitoring a selected tier in a multi-tier computing environment, said selected tier is coupled with at least one adjacent tier of said multi-tier computing environment, said apparatus comprising:

20 a context agent associated with said selected tier, said context agent being coupled with at least one other context agent, each of said at least one other context agent being associated with a respective one of said at least one adjacent tier of said computing environment; and

a dynamic tier extension, said dynamic tier extension coupled with said
25 context agent and with at least a request entry port and a request exit port of said

77964-10

selected tier, said dynamic tier extension at least monitoring request traffic passing through said selected tier, said monitored request traffic including at least one entering request received at said request entry port and at least one respective exiting request at said request exit port, said selected tier providing said at least one
 5 respective exiting request to a receiving tier of said at least one adjacent tier, said dynamic tier extension capturing a request context of said at least one entering request, said dynamic tier extension assigning a unique transaction identifier to said at least one entering request for uniquely identifying each request of a transaction to which said at least one entering request belongs, said dynamic tier extension
 10 providing request information and context information to said context agent, said request information at least identifying each request included in said monitored request traffic, said context information relating to said captured request context of said at least one entering request, said context information including at least said unique transaction identifier,

15 wherein said context agent associates said at least one entering request with said at least one exiting request, and

wherein said context agent provides context information relating to the request context of said at least one respective exiting request directly to another context agent associated with said receiving tier of said at least one adjacent tier.

20 24. The apparatus according to Claim 23, wherein said at least one entering request is a user request received from a user application.

25. The apparatus according to Claim 23, wherein said context information includes at least a request identifier, a request class, and context related data.

26. The apparatus according to Claim 23, wherein said context agent
 25 classifies said at least one entering request into a request class in accordance with said context information.

77964-10

27. The apparatus according to Claim 26, wherein said context agent contains a set of tier-specific service class policies, each of said tier-specific service class policies mapping a service class to a request class for said selected tier,

wherein an active service class policy contains the request class to
5 service class mapping that is currently in effect; and

wherein said context agent assigns a service class to said at least one entering request based on said request class and said active service class policy in said set of tier-specific service class policies.

28. The apparatus according to Claim 26, wherein said context agent
10 contains a set of tier-specific service class policies, each of said tier-specific service class policies mapping a service class to a request class for said selected tier,

wherein an active service class policy contains the request class to service class mapping that is currently in effect; and

wherein said context agent assigns a service class to said at least one
15 entering request different than the service class designated by said request class and said active service class policy in said set of tier-specific service class policies.

29. The apparatus according to either of Claims 27 or 28, wherein said service class includes a priority to be assigned to said at least one entering request.

30. The apparatus according to either of Claims 27 or 28, wherein said
20 service class includes a central processing unit percentage to be allocated to said at least one entering request.

31. The apparatus according to either of Claims 27 or 28, wherein said service class includes a memory to be allocated to said at least one entering request.

77964-10

32. The apparatus according to either of Claims 27 or 28, wherein said service class includes a priority in allocating and accessing input and output devices to said at least one entering request.

33. The apparatus according to either of Claims 27 or 28, wherein said
5 service class policies are periodically updated.

34. The apparatus according to either of Claims 27 or 28, wherein said service class policies are defined by a user of said multi-tier computing environment.

35. The apparatus according to either of Claims 27 or 28, wherein said context agent adds said service class, assigned by said context agent, to said
10 request context.

36. The apparatus according to Claim 23, wherein said dynamic tier extension is further coupled with a tier control port of said selected tier.

37. The apparatus according to Claim 36, wherein said context agent further alters a priority level of a processing enclave executing said at least one
15 entering request.

38. The apparatus according to Claim 36, wherein said context agent further alters the type of execution of a processing enclave executing said at least one entering request.

39. The apparatus according to either of Claims 23 or 36, wherein said
20 context agent further adjusts the order of said at least one entering request on a queue at said request entry port.

40. The apparatus according to either of Claims 23 or 36, wherein said context agent further instructs said selected tier to allocate computational resources to process said at least one entering request.

25 41. The apparatus according to either of Claims 23 or 36, wherein said context agent further instructs said selected tier to deny computational resources to

77964-10

process said at least one entering request.

42. The apparatus according to Claim 23, wherein said dynamic tier extension assigns a transaction identifier to each request included in said monitored request traffic.

5 43. The apparatus according to Claim 23, wherein said dynamic tier extension obtains a request identifier of each request included in said monitored request traffic.

44. The apparatus according to Claim 23, wherein said context agent comprises:

10 an agent communication interface, relaying messages between said dynamic tier extension and said context agent;

an agent messaging module, notifying said another context agent that said at least one exiting request has been sent to said another adjacent tier, said agent messaging module further communicating with said dynamic tier extension
15 using messages;

a context table manager, maintaining a reference table associating at least one unit of work running on said selected tier with said request context of said at least one entering request to which said unit of work has been allocated, said reference table being a context table, said context table further storing information
20 associated with said at least one entering request, said context table manager further identifying at least one entering request based on information stored in said context table;

a classification manager, performing classification of said at least one entering request based on said received context information, said classification
25 manager further identifying said at least one entering request as part of a transaction; and

77964-10

a policy manager, assigning a service class policy to said at least one entering request according to an active service class policy and said classification of said at least one entering request.

45. The apparatus according to Claim 44, wherein said context agent
5 further comprises:

an agent log event handler, performing event logging, internal tracing and messaging, and detection of operational irregularities within said context agent.

46. The apparatus according to Claim 44, wherein said context agent
further comprises:

10 an agent administrator, archiving information, sending historical data to a context network management server (CNMS) and receiving active service class policies and additional data from said CNMS.

47. The apparatus according to Claim 36, wherein said dynamic tier
extension comprises:

15 a dynamic tier extension daemon, performing asynchronous processing associated with reported events where said at least one entering request is not detained;

a dynamic tier extension messaging module, communicating with said context agent using messages; and

20 a dynamic tier extension communication interface, relaying messages between said dynamic tier extension and said context agent.

48. The apparatus according to Claim 47, wherein said dynamic tier
extension further comprises:

77964-10

a dynamic tier extension administrator, enabling said dynamic tier extension to receive messages relating to how said at least one entering request should be processed.

49. The apparatus according to Claim 47, wherein said dynamic tier
5 extension further comprises:

a dynamic tier extension log event handler, logging tracing information concerning operation of said dynamic tier extension and logging alerts raised by said dynamic tier extension.

50. The apparatus according to Claim 23, wherein said context agent
10 resides on the same host machine as said selected tier.

51. The apparatus according to Claim 36, wherein said dynamic tier extension collects performance, availability, and error metrics from said selected tier.

52. In a multi-tier computing environment including a plurality of tiers, a system for application performance management, associated with at least two
15 selected monitored tiers of said plurality of tiers, the system comprising:

for each of at least two of said monitored tiers, a respective dynamic tier extension being coupled with at least a request entry port of said monitored tier, said dynamic tier extension at least monitoring request traffic passing through said respective monitored tier, said monitored request traffic including at least one
20 entering request received at said request entry port; and

for each of said at least two of said monitored tiers, a respective context agent, coupled with said respective dynamic tier extension, said context agent being coupled with other said context agents associated with tiers that are directly coupled with said respective tier, said respective context agent receiving request information
25 from said respective dynamic tier extension, said request information at least identifying each request included in said monitored request traffic, said respective

77964-10

context agent further receives context information relating to the request context of said at least one entering request directly from another context agent associated with a respective one of said tiers that are directly coupled with said respective tier, said received context information including at least a unique transaction identifier for

5 uniquely identifying each request of a transaction to which said at least one entering request belongs, for each said at least one entering request, said respective context agent associates said context information with said request information, respectively.

53. In a multi-tier computing environment including a plurality of tiers, a system for application performance management, associated with at least two
10 selected monitored tiers of said plurality of tiers, the system comprising:

for each of at least two of said monitored tiers, a respective dynamic tier extension being coupled with at least a request entry port and a request exit port of said monitored tier, said dynamic tier extension at least monitoring request traffic passing through said respective monitored tier, said monitored request traffic
15 including at least one entering request received at said request entry port and at least one respective exiting request at said request exit port, said selected tier providing said at least one respective exiting request to an adjacent respective monitored tier, said dynamic tier extension capturing a request context of said at least one entering request, said dynamic tier extension assigning a unique transaction identifier to said
20 at least one entering request for uniquely identifying each request of a transaction to which said at least one entering request belongs; and

for each of said at least two of said monitored tiers, a respective context agent, coupled with said respective dynamic tier extension, said context agent being coupled with other said context agents associated with tiers that are directly coupled
25 with said respective tier, said respective context agent receiving request information from said respective dynamic tier extension, said request information at least identifying each request included in said monitored request traffic, said respective context agent further receives context information relating to the request context of said at least one entering request from said respective dynamic tier extension, for

77964-10

each said at least one entering request, said respective context agent associates said at least one entering request with said at least one exiting request, and said respective context agent provides context information relating to the request context of said at least one respective exiting request directly to another context agent associated with said adjacent respective monitored tier, said context information including at least said unique transaction identifier.

54. The system according to either of Claims 52 or 53, further comprising:

a context network management server (CNMS) coupled with said context agents, said CNMS at least collecting and analyzing performance data received from said context agents.

55. The system according to either of Claims 52 or 53, wherein said at least one entering request is a user request received from a user application.

56. The system according to Claim 52, wherein each of said dynamic tier extensions is further coupled with at least a request exit port of said respective tier, and wherein said monitored request traffic further includes at least one exiting request received at said request exit port, said selected tier providing said at least one exiting request to another adjacent tier of said computing environment,

wherein said respective context agent associates said at least one exiting request with said at least one entering request; and

sends context information relating to the request context of said at least one exiting request to a further context agent associated with said another adjacent tier.

57. The system according to Claim 54, further including a policy database coupled with said CNMS, said policy database storing at least service class policies.

58. The system according to Claim 57, further including a supervisor workstation coupled with said CNMS and with said policy database.

77964-10

59. The system according to Claim 54, wherein said CNMS further performs activity profiling, creating an activity profile including integrated transaction, tier, and system level performance metrics and statistical analysis.

60. In a multi-tier computing environment including a plurality of tiers, an
5 application performance management method for monitoring a selected monitored tier of said plurality of tiers, the method comprising the procedures of:

receiving at a context agent associated with said selected monitored tier context information relating to a request context of at least one entering request from another context agent associated with another tier which sent said at least one
10 entering request to said selected monitored tier, said context information including at least a request identifier and a unique transaction identifier for uniquely identifying each request of a transaction to which said at least one entering request belongs;

monitoring request traffic passing through said monitored tier, said monitored request traffic including at least said at least one entering request;

15 identifying said at least one entering request in accordance with said request identifier; and

associating said at least one entering request with a transaction to which said at least one entering request belongs in accordance with said transaction identifier.

20 61. In a multi-tier computing environment including a plurality of tiers, an application performance management method for monitoring a selected monitored tier of said plurality of tiers, the method comprising the procedures of:

monitoring request traffic passing through said monitored tier, said monitored request traffic including at least one entering request and at least one
25 respective exiting request, said selected tier providing said at least one respective exiting request to an adjacent tier;

77964-10

determining context information relating to the request context of said at least one entering request, said context information including at least a unique transaction identifier for uniquely identifying each request of a transaction to which said at least one entering request belongs;

5 identifying each request included in said monitored request traffic;

associating said at least one entering request with said at least one exiting request; and

sending context information relating to the request context of said at least one exiting request directly to another context agent associated with said
10 adjacent tier.




FIG. 1
PRIOR ART




FIG. 2





FIG. 3

**FIG. 4**









FIG. 5


**FIG. 6**

**FIG. 7**

**FIG. 8**

**FIG. 9**

**FIG. 10**

**FIG. 11**

