

(43) Date of A Publication **08.12.1999**

(21) Application No **9905212.8**

(22) Date of Filing **05.03.1999**

(30) Priority Data

(31) **09089848** (32) **03.06.1998** (33) **US**
(31) **60077157** (32) **06.03.1998**

(71) Applicant(s)

Mosaix Inc
(Incorporated in USA - Washington)
6464-185th Avenue NE, Redmond,
Washington 98052-6736, United States of America

(72) Inventor(s)

Norman J Donaghue Jr

(74) Agent and/or Address for Service

Withers & Rogers
Goldings House, 2 Hays Lane, LONDON, SE1 2HW,
United Kingdom

(51) INT CL⁶

G06F 9/46 17/60

(52) UK CL (Edition Q)

G4A AFN

(56) Documents Cited

US 5504894 A **US 5473773 A**

(58) Field of Search

UK CL (Edition Q) G4A AFN
INT CL⁶ G06F 9/46 17/60
Online: WPI, EPODOC, JAPIO

(54) Abstract Title

Prioritized transaction server allocation

(57) A method for distributing transaction servers to process transactions of a plurality of different types organised into a plurality of priority levels ordered from highest to lowest, each transaction type having a service level goal, comprises the steps of identifying available transaction servers, determining, from the highest to the lowest priority level, whether assigning a server to a transaction type of the priority level is necessary to meet the service level goal for that type, and if this is necessary, assigning the server to the transaction, and after the server has completed processing of the transaction, reidentifying the server as being available, such that an available server cannot be assigned to a transaction of a lower priority level when the service level goal of the transaction type of the higher priority level is not satisfied. The method may be applied to processing telephone call transactions in a call centre.

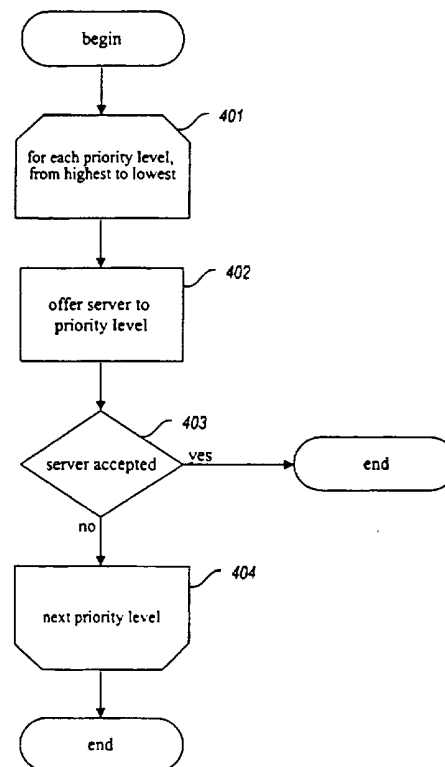


Fig. 4

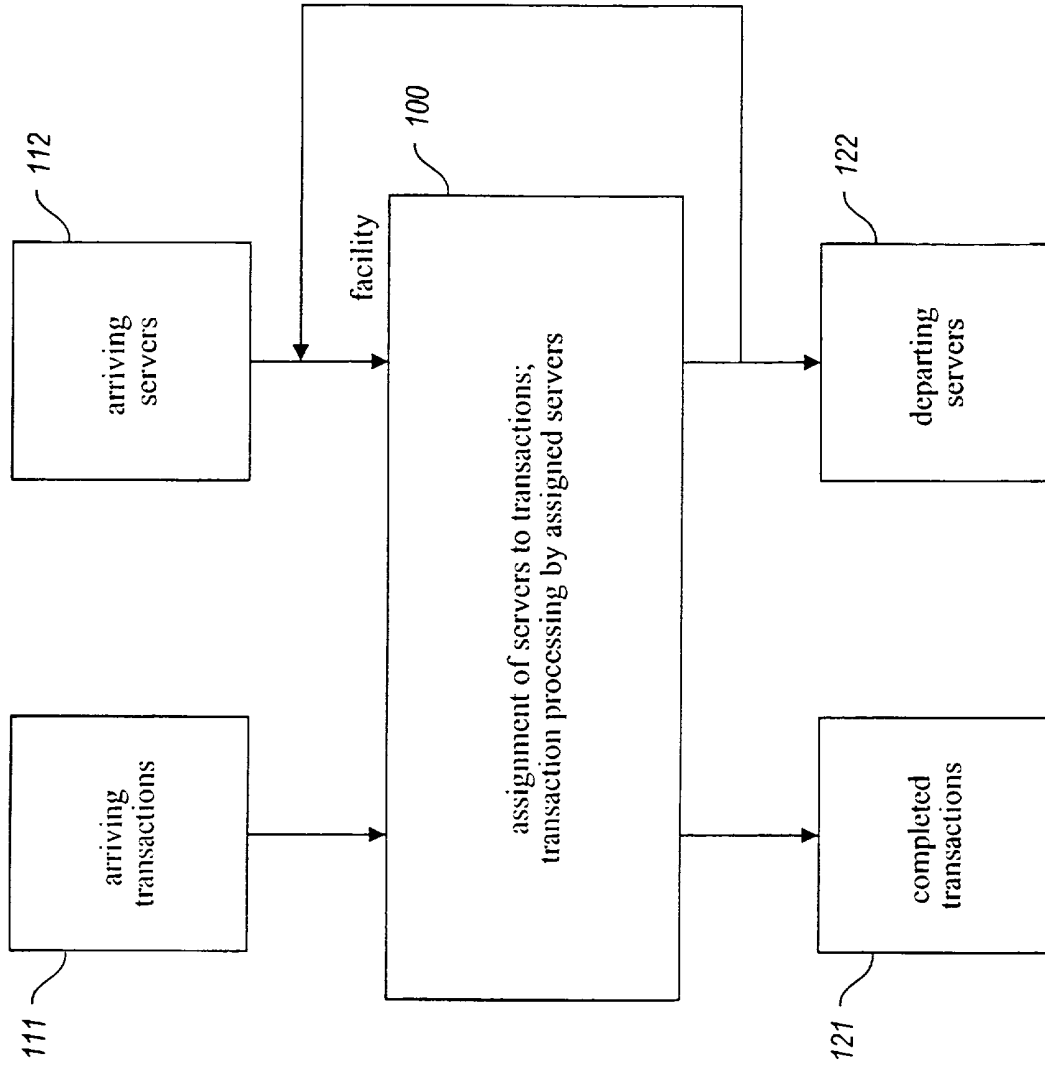

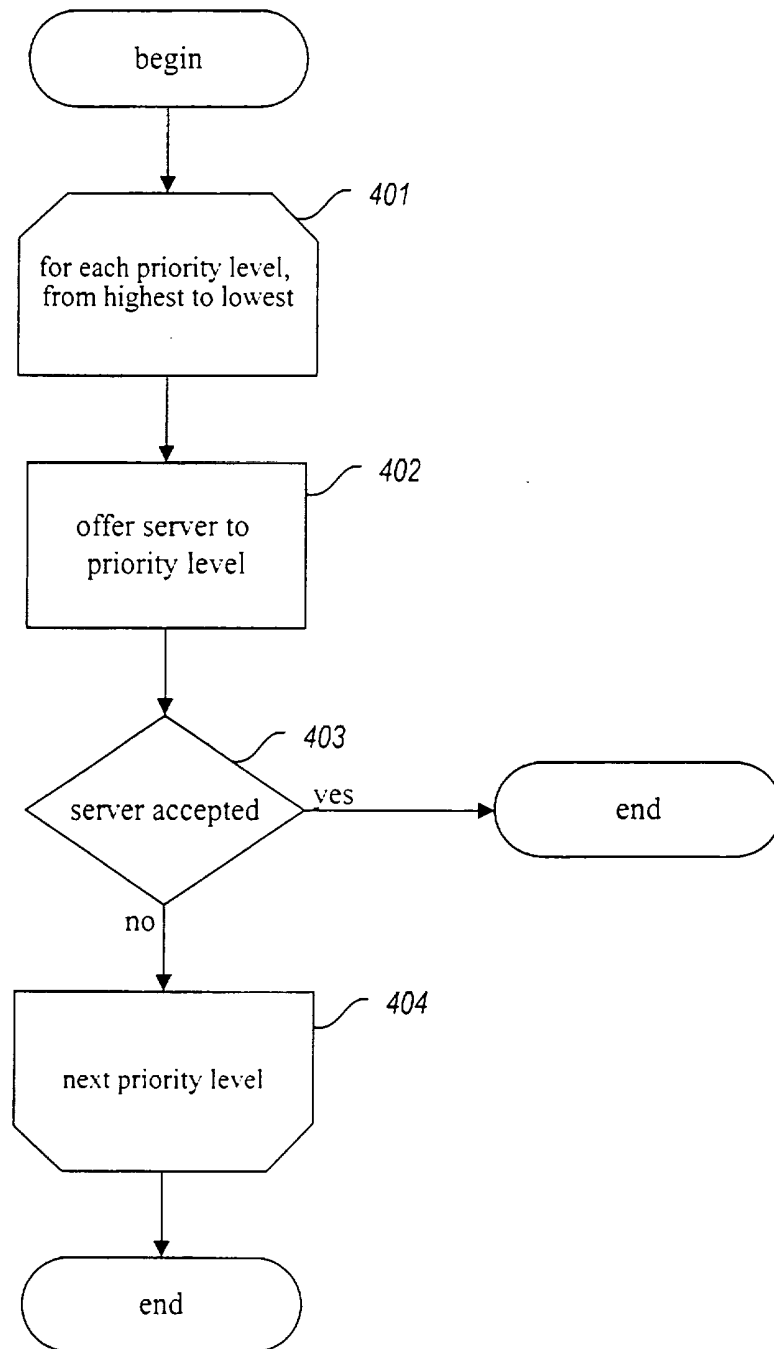


Fig. 1

application number	priority level	transaction type/application	service level goal		service level goal type
			performance criterion	minimum frequency	
1	1	incoming platinum calls	assign within 15 seconds	90% of transactions	transaction assignment
2	2	incoming gold calls	assign within 60 seconds	90% of transactions	transaction assignment
3	2	credit check	complete within 4 days	80% of transactions	transaction completion
4	3	outgoing card offer calls	complete 200 transactions per day	75% of days	transaction completion
5	3	outgoing brochures	complete within 4 days	70% of transactions	transaction completion
6	3	responses to incoming letters	complete within 7 days	75% of transactions	transaction completion
7	100	at-desk break	accept all offered agents		none

Fig. 2

*Fig. 3*

*Fig. 4*

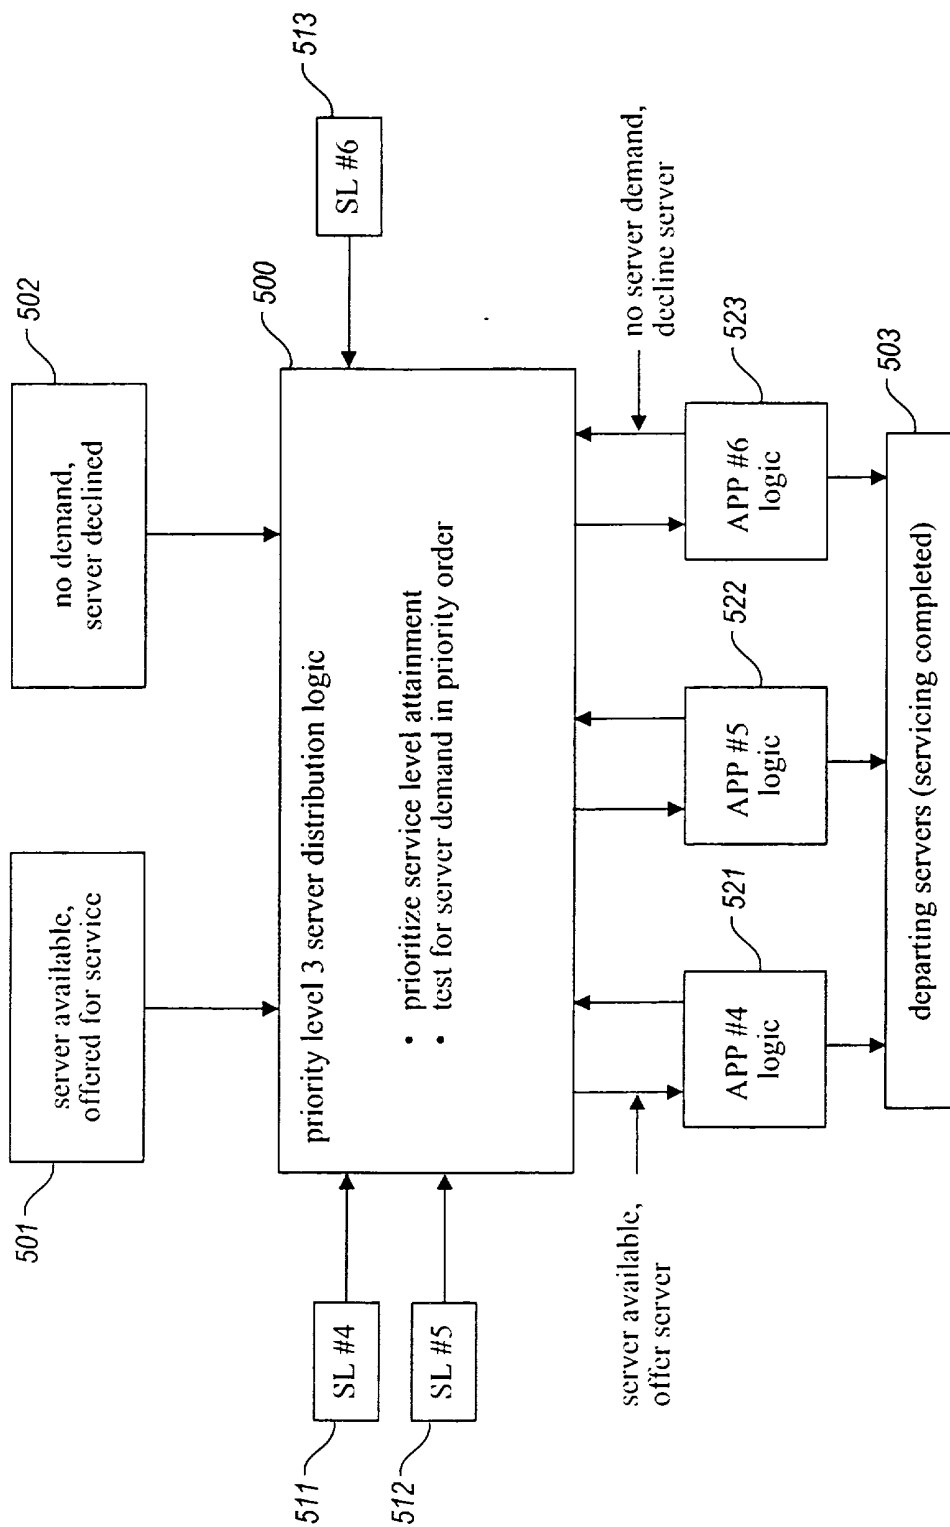

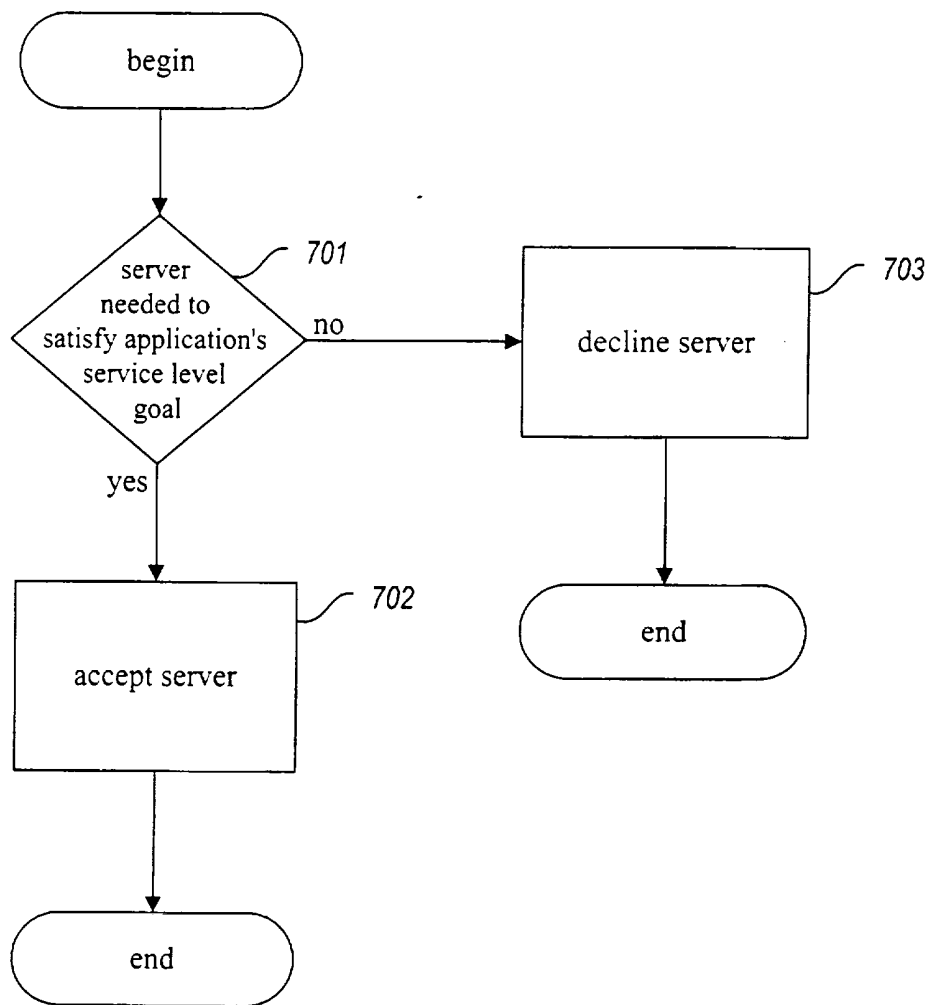


Fig. 5

*Fig. 6*

*Fig. 7*

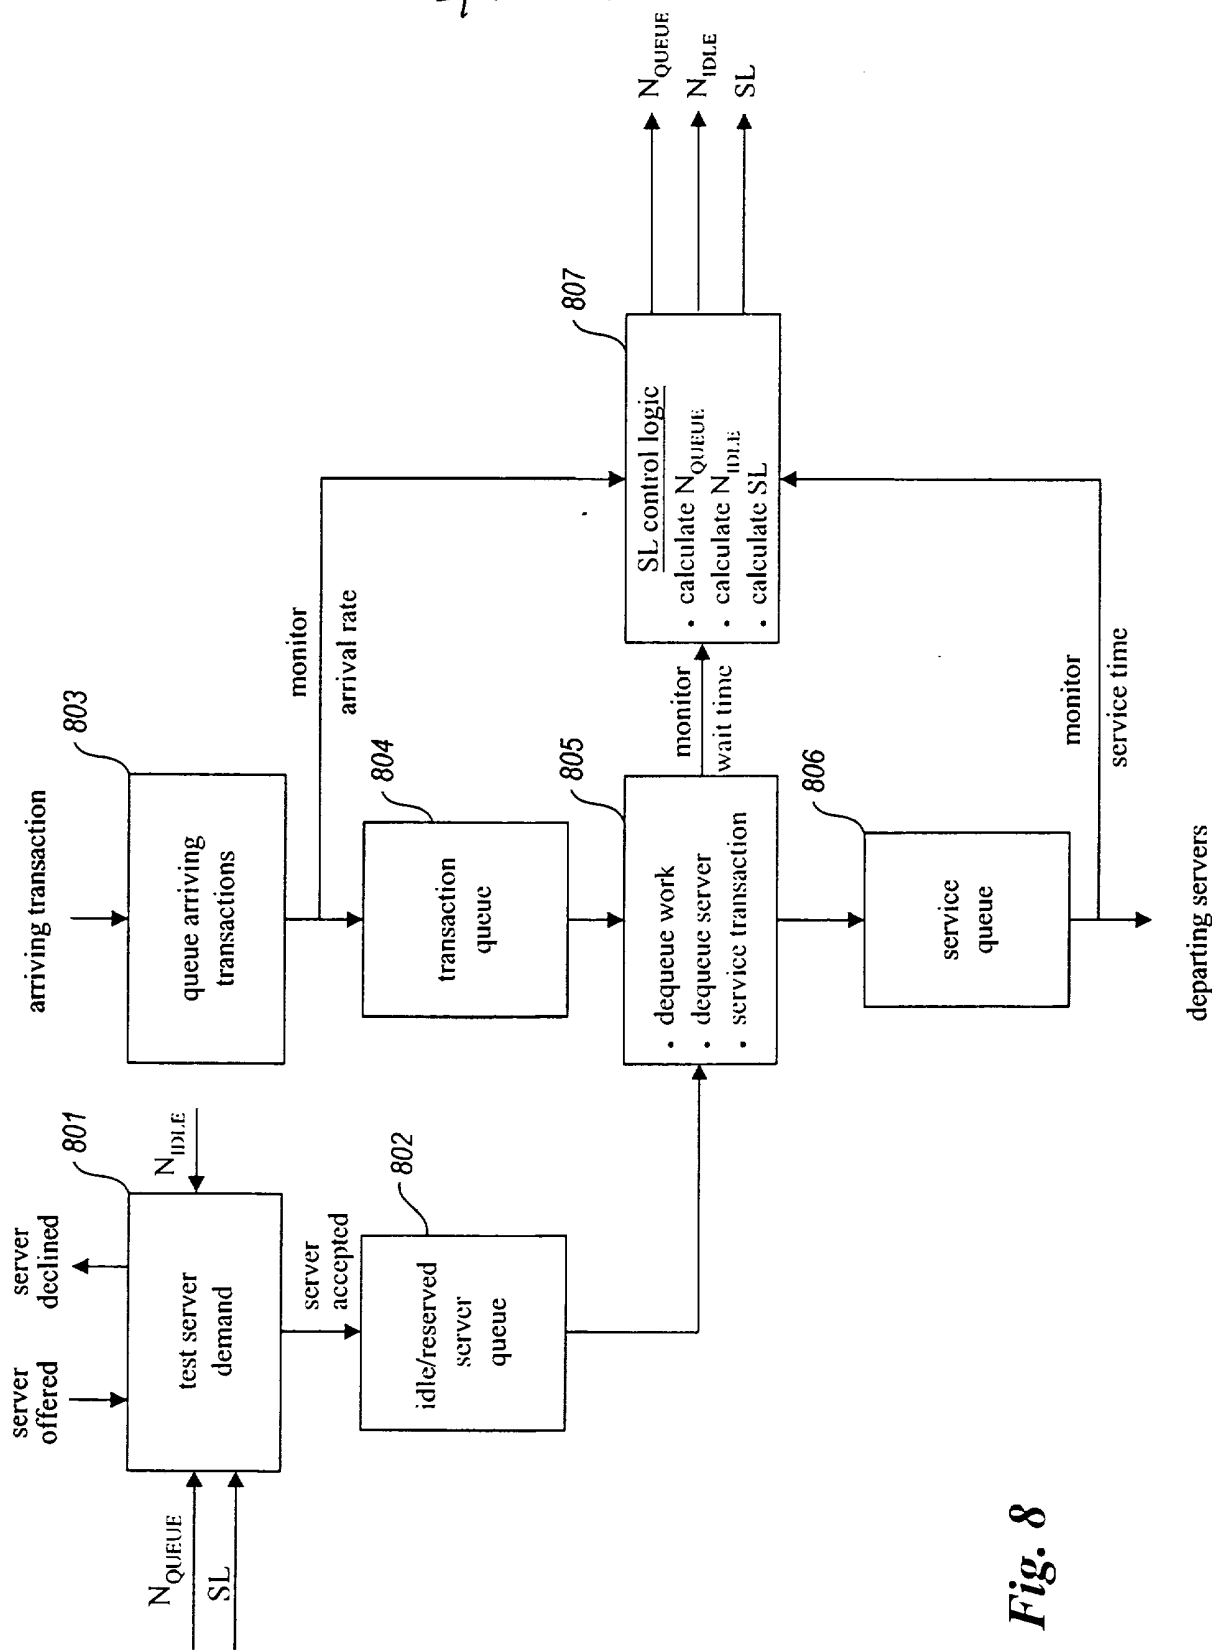


Fig. 8




Fig. 9

PRIORITIZED TRANSACTION SERVER ALLOCATION

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority from U.S. Provisional Application No. 60/077,157, filed March 6, 1998, which application is incorporated by reference in
5 its entirety.

TECHNICAL FIELD

The present invention relates generally to the field of allocating transaction server resources and, in particular, to allocating transaction servers in telephonic call centers.

10 BACKGROUND OF THE INVENTION

A call center is a facility for receiving and/or placing large volumes of telephone calls. The calls are received and placed by call servers, such as human agents or automatic response units. Such call centers typically include an automatic call distributor for assigning each incoming call to a call server. Call centers also typically
15 include a predictive dialer for placing outgoing calls that, if answered, are each assigned to a call server.

In some call centers, agents are also assigned to non-call tasks, such as filling out forms, writing letters, or running credit checks. Call centers in which agents are also assigned to non-call tasks are also called service centers. Assignment of these
20 non-call tasks is generally accomplished manually or handled by a workflow system that is distinct from the automatic call distributor or predictive dialer that assigns phone calls to agents.

Conventionally, the automatic call distributor, the predictive dialer, and the workflow system operate independently to assign calls and non-call tasks,
25 respectively, to agents. To accommodate the independence between these two components, a call center supervisor generally manually divides agent resources between calls and non-call tasks in a static, predetermined manner. For example, a first

group of agents may be dedicated to conducting calls, while a second, distinct group of agents is dedicated to performing non-call tasks. As another example, certain agents may be dedicated to conducting calls every morning, and dedicated to performing non-call tasks every afternoon.

5 Such call centers usually succeed in making a certain amount of progress in handling both calls and non-call tasks. Depending on the aptness of the manual division of agent resources and the particular pattern of calls and non-call tasks experienced, however, unacceptable levels of service are often obtained for either calls, non-call tasks, or both. In cases in which it is essential to ensure high levels of service,
10 it is often necessary for call center supervisors to overstaff the pool of agents, at significant cost.

 Based on these disadvantages of conventional call centers, which are shared by similar transaction processing systems of several other types, a facility for more effectively managing call server resources and other transaction server resources
15 would have significant utility.

SUMMARY OF THE INVENTION

 The present invention provides a facility for allocating transaction servers ("servers") in a transaction processing system in a manner that facilitates differentiation between transactions of different priorities and balancing the
20 performance of transactions of the same priority. Within the constraints imposed by the availability of servers, the facility facilitates maintaining pre-defined service level goals. For example, the facility may allocate telephone call processing servers, such as human operators/agents and automatic response units, for processing telephone calls. The facility can also allocate servers to process virtually any other type of transaction,
25 including, but not limited to, form processing; response to world wide web hits ("web hits"); letter, fax, or email generation or response; account adjustment; and various forms of network and Internet communication. The facility can allocate servers between several different types of transactions. Transactions may be initiated within

the organization using the facility, by external customers or clients of the organization, or both.

In a preferred embodiment, the facility allocates servers among different applications each of which processes transactions of a different type or of the same type but with a different service level goal. For example, a first application may process
5 incoming telephone call transactions for platinum credit card holders, a second application may process outgoing telephone call transactions for new credit card offers, and a third application may process credit check transactions, such as those received via the world wide web or in the mail, and a fourth application that responds to web hits
10 with outgoing phone calls.

The facility preferably enables a service level goal to be specified for each application, indicating the level of performance sought in processing transactions of the application's type. Service level goals are generally each comprised of two components: a performance criterion specifying an amount of time by which some
15 action relating to the transaction is to be completed, and a minimum frequency indicating the percentage of transactions for which the criterion is to be satisfied. As an example, one service level goal may specify a criterion of assigning transaction within 15 seconds of their receipt, and a minimum frequency that indicates that 90% of the transactions are to be processed in accordance with the performance criterion. Performance criteria
20 ("criteria") may be of various types, including transaction *assignment* criteria and transaction *completion* criteria. Transaction assignment criteria relate to the amount of time that it takes for a server to be assigned to transactions after they arrive. A transaction assignment criterion may indicate, for example, that incoming telephone call transactions for platinum credit card holders are to be assigned to a server within 15
25 seconds of their arrival. Transaction assignment criteria are particularly well adapted to transactions initiated outside the organization using a real-time medium, such as an incoming telephone call or web hit—especially those transactions initiated by a human client. Such transactions are regarded as "volatile," as they may be lost if not quickly assigned to a server. For example, if an incoming telephone call is not answered within
30 a few seconds after it is initiated, the caller may hang up and end the call.

Transaction completion criteria, on the other hand, relate to the amount of time that it takes to complete the processing of transactions. A transaction completion criterion may indicate, for example, that credit check transactions are to be completed within 4 days of their arrival, or that outgoing telephone call transactions for new credit card offers are to be completed at the rate of 200 per day. Transaction completion criteria are particularly well adapted to transactions initiated outside the organization using a non-real-time medium, such as incoming letters, or transactions that are initiated from inside the organization, such as outgoing telephone calls or are generated as a follow on to a transaction assignment transaction, such as a telephone call requesting a loan. Servicing the telephone call involves a transaction assignment criterion (such as answering all customer telephone calls within 20 seconds) but generating the loan application that results involves satisfying a transaction completion criterion (such as processing all loan applications within 4 business days).

In addition to specifying a performance criterion, a service level goal may also indicate a minimum percentage of times, or "minimum frequency," that the performance criterion is to be satisfied. For example, a transaction assignment service level goal may indicate that *at least 90%* of incoming telephone call transactions for platinum credit card holders are to be assigned to a server within 15 seconds of their arrival.

In addition to a service level goal, each application preferably also has a priority level, indicating the importance of satisfying the service level goal of that application relative to the importance of satisfying the service level goals of the other applications. For example, incoming telephone call transactions for platinum credit card holders may have a priority level of 1 (indicating the highest priority), incoming telephone call transactions for gold credit card holders may have a priority level of 2, and outgoing telephone call transactions for new credit card offers may have a priority level of 3. Multiple applications may have the same priority level.

In contrast to conventional transaction processing systems that attempt to assign each newly-arrived transaction to a server, the facility of the present invention assigns each newly-available server to a transaction. Such assignment takes place each

time a server becomes available, either when an existing server completes the processing of the last transaction that it was assigned to process, or when a new server arrives. The assignment process involves "offering" use of the server for use by each priority level of applications in sequence, from the highest priority level to the lowest priority level, until a priority level accepts use of the server. To determine whether to accept the offered server, each priority level in turn offers the server to each application of that priority level, until an application accepts use of the server. If any application at that priority level accepts use of the server, then the priority level accepts the use of the server. On the other hand, if none of the applications at that priority level accepts use of the server, then the priority level declines the use of the server, and the facility offers use of the server to the next lower priority level. By offering each available server in accordance with the priority levels, the facility ensures that at least applications in the highest priority levels will successfully reach their service level goals. Where adequate numbers of servers are available, however, the facility permits applications in all priority levels to reach their service level goals.

Within a single priority level, the facility offers use of the server to the applications of the priority level in a certain sequence. In cases in which the server is needed by more than one application in the current priority level, this sequence determines which of these applications is able to use the server. The facility therefore establishes the sequence in a manner that is fair to all of the applications of the priority level. As one example, the facility may use a round-robin approach for establishing this order, beginning with an initial random order and, for each offer, rotating the order one position in a consistent direction. In this way, each application of the priority level has equal opportunity to accept servers offered to the priority level. As another example, the facility may assess the success of each application at reaching its service level goals, and establish a sequence that favors the applications that are the furthest from reaching their service level goals. In this way, the facility exerts a balancing force on the applications of the priority level, enabling underperforming applications to receive extra servers that can be used to improve the performance of these applications.

When the facility offers a server to a particular application, the application preferably determines whether it should accept the offered server based on whether the offered server is needed by the application to meet its service level goal. The facility may make this determination either on a "closed-loop" basis, in which the
5 facility uses the current achieved service level of the application to make the determination, or on an "open-loop" basis, in which the facility uses information other than the current achieved service level of the application.

For the closed loop basis, the facility compares the current achieved service level with the pre-defined service level goal to make the determination. If the
10 offered server is not required for achievement of the predetermined service level goals, then the server is not accepted. If the offered server is required to achieve the predetermined service level goal, then the server is accepted.

For the open-loop basis, the facility may use a preestablished model, formula, or table to determine whether an offered server is needed by the application to
15 meet its service level goal using current operating parameters of the application, such as transaction arrival rate, number of uninitiated transactions, wait times of uninitiated transactions, transaction service completion rate, and transaction service completion time.

When the application accepts the offered server, if any transactions are
20 waiting for processing, the accepted server is assigned to process one of the waiting transactions. On the other hand, if no transactions are waiting when the application accepts the server, the accepted server is maintained in an idle/reserved server queue until a new transaction arrives. When a new transaction subsequently arrives, a server in the idle/reserved server queue is assigned to process the new transaction. By
25 maintaining one or more servers in its idle/reserved server queue, that is, by accepting a server or servers even when its transaction queue is empty, an application can ensure that a server will be immediately assigned to the next arriving transaction. In this way, the application is able to exert control over its achieved service level.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a conceptual diagram showing the relationship between the facility, transactions, and servers.

Figure 2 is a table diagram identifying a number of sample transaction
5 types.

Figure 3 is a conceptual diagram showing the facility's approach to assigning available servers to the priority levels identified in Figure 3.

Figure 4 is a flow diagram showing the steps preferably performed by the facility in order to assign an available server to one of the priority levels.

10 Figure 5 is a conceptual diagram showing the facility's approach to assigning available servers to the sample applications identified in Figure 2 within a particular priority level.

Figure 6 is a flow diagram showing the steps preferably performed by the facility in order to offer an available server to a particular priority level, as described
15 above for step 402 of Figure 4.

Figure 7 is a flow diagram showing the steps preferably performed by the facility in order to accept or decline an available server offered to a particular application, as described above for step 603 of Figure 6.

20 Figure 8 is a workflow diagram showing the activity performed by the facility with respect to a single application.

Figure 9 is a high-level block diagram showing an illustrative computer system upon which the facility of the present invention may operate.

DETAILED DESCRIPTION OF THE INVENTION

25 The present invention provides a facility for allocating transaction servers ("servers") in a transaction processing system in a manner that facilitates differentiation between transactions of different priorities and balancing the performance of transactions of the same priority. Within the constraints imposed by the availability of servers, the facility facilitates maintaining pre-defined service level goals. For example, the facility may allocate telephone call processing servers, such as human

operators/agents and automatic response units, for processing telephone calls. The facility can also allocate servers to process virtually any other type of transaction, including, but not limited to, form processing; response to world wide web hits ("web hits"); letter, fax, or email generation or response; account adjustment; and various forms of network and Internet communication. The facility can allocate servers between several different types of transactions. Transactions may be initiated within the organization using the facility, by external customers or clients of the organization, or both.

In a preferred embodiment, the facility allocates servers among different applications, each of which processes transactions of a different type or of the same type but with a different service level goal. For example, a first application may process incoming telephone call transactions for platinum credit card holders, a second application may process outgoing telephone call transactions for new credit card offers, and a third application may process credit check transactions, such as those received via the world wide web or in the mail, and a fourth application that responds to web hits with outgoing phone calls.

The facility preferably enables a service level goal to be specified for each application, indicating the level of performance sought in processing transactions of the application's type. Service level goals are generally each comprised of two components: a performance criterion specifying an amount of time by which some action relating to the transaction is to be completed, and a minimum frequency indicating the percentage of transactions for which the criterion is to be satisfied. As an example, one service level goal may specify a criterion of assigning transaction within 15 seconds of their receipt, and a minimum frequency that indicates that 90% of the transactions are to be processed in accordance with the performance criterion. Performance criteria ("criteria") may be of various types, including transaction *assignment* criteria and transaction *completion* criteria. Transaction assignment service level goals relate to the amount of time that it takes for a server to be assigned to transactions after they arrive. A transaction assignment criterion may indicate, for example, that incoming telephone call transactions for platinum credit card holders are to be assigned to a server within 15

seconds of their arrival. Transaction assignment criteria are particularly well adapted to transactions initiated outside the organization using a real-time medium, such as an incoming telephone call or web hit—especially those transactions initiated by a human client. Such transactions are regarded as “volatile,” as they may be lost if not quickly
 5 assigned to a server. For example, if an incoming telephone call is not answered within a few seconds after it is initiated, the caller may hang up and end the call.

Transaction completion criteria, on the other hand, relate to the amount of time that it takes to complete the processing of transactions. A transaction completion criterion may indicate, for example, that credit check transactions are to be
 10 completed within 4 days of their arrival, or that outgoing telephone call transactions for new credit card offers are to be completed at the rate of 200 per day. Transaction completion criteria are particularly well adapted to transactions initiated outside the organization using a non-real-time medium, such as incoming letters, or transactions that are initiated from inside the organization, such as outgoing telephone calls or are
 15 generated as a follow on to a transaction assignment transaction, such as a telephone call requesting a loan. Servicing the telephone call involves a transaction assignment criterion (such as answering all customer telephone calls within 20 seconds) but generating the loan application that results involves satisfying a transaction completion criterion (such as processing all loan applications within 4 business days).

20 In addition to specifying a performance criterion, service level goals may also indicate a minimum percentage of times, or “minimum frequency,” that the performance criterion is to be satisfied. For example, a transaction assignment service level goal may indicate that *at least 90%* of incoming telephone call transactions for platinum credit card holders are to be assigned to a server within 15 seconds of their
 25 arrival.

In addition to a service level goal, each application preferably also has a priority level, indicating the importance of satisfying the service level goal of that application relative to the importance of satisfying the service level goals of the other applications. For example, incoming telephone call transactions for platinum credit
 30 card holders may have a priority level of 1 (indicating the highest priority), incoming

telephone call transactions for gold credit card holders may have a priority level of 2, and outgoing telephone call transactions for new credit card offers may have a priority level of 3. Multiple applications may have the same priority level.

In contrast to conventional transaction processing systems that attempt to
5 assign each newly-arrived transaction to a server, the facility of the present invention assigns each newly-available server to a transaction. Such assignment takes place each time a server becomes available, either when an existing server completes the processing of the last transaction that it was assigned to process, or when a new server arrives. The assignment process involves "offering" use of the server for use by each
10 priority level of applications in sequence, from the highest priority level to the lowest priority level, until a priority level accepts use of the server. To determine whether to accept the offered server, each priority level in turn offers the server to each application of that priority level, until an application accepts use of the server. If any application at that priority level accepts use of the server, then the priority level accepts the use of the
15 server. On the other hand, if none of the applications at that priority level accepts use of the server, then the priority level declines the use of the server, and the facility offers use of the server to the next lower priority level. By offering each available server in accordance with the priority levels, the facility ensures that at least applications in the highest priority levels will successfully reach their service level goals. Where adequate
20 numbers of servers are available, however, the facility permits applications in all priority levels to reach their service level goals.

Within a single priority level, the facility offers use of the server to the applications of the priority level in a certain sequence. In cases in which the server is needed by more than one application in the current priority level, this sequence
25 determines which of these applications is able to use the server. The facility therefore establishes the sequence in a manner that is fair to all of the applications of the priority level. As one example, the facility may use a round-robin approach for establishing this order, beginning with an initial random order and, for each offer, rotating the order one position in a consistent direction. In this way, each application of the priority level has
30 equal opportunity to accept servers offered to the priority level. As another example,

the facility may assess the success of each application at reaching its service level goals, and establish a sequence that favors the applications that are the furthest from reaching their service level goals. In this way, the facility exerts a balancing force on the applications of the priority level, enabling underperforming applications to receive extra
5 servers that can be used to improve the performance of these applications.

When the facility offers a server to a particular application, the application preferably determines whether it should accept the offered server based on whether the offered server is needed by the application to meet its service level goal. The facility may make this determination either on a "closed-loop" basis, in which the
10 facility uses the current achieved service level of the application to make the determination, or on an "open-loop" basis, in which the facility uses information other than the current achieved service level of the application.

For the closed loop basis, the facility compares the current achieved service level with the pre-determined service level goal to make the determination. If
15 the offered server is not required for achievement of the predetermined service level goal, then the server is not accepted. If the offered server is required to achieve the pre-determined service level goal, then the server is accepted.

For the open-loop basis, the facility may use a preestablished model, formula, or table to determine whether an offered server is needed by the application to
20 meet its service level goal using current operating parameters of the application, such as transaction arrival rate, number of uninitiated transactions, wait times of uninitiated transactions, transaction service completion rate, and transaction service completion time.

When the application accepts the offered server, if any transactions are
25 waiting for processing, the accepted server is assigned to process one of the waiting transactions. On the other hand, if no transactions are waiting when the application accepts the server, the accepted server is maintained in a reserved/idle server queue until a new transaction arrives. When a new transaction subsequently arrives, a server in the reserved/idle server queue is assigned to process the new transaction. By maintaining
30 one or more servers in its reserved/idle server queue, that is, by accepting a server or

servers even when its transaction queue is empty, an application can ensure that a server will be immediately assigned to the next arriving transaction. In this way, the application is able to exert control over its achieved service level.

Figure 1 is a conceptual diagram showing the relationship between the facility, transactions, and servers. It can be seen from Figure 1 that the facility 100 receives arriving transactions 111, such as telephone calls or web hits, and arriving servers 112, such as servers like human agents, automatic telephone response units, and/or automated web servers. Servers may arrive because they have just become available ("logged on" or "gone online"), or because they have completed servicing a transaction and are available for reassignment. The facility 100 assigns each arriving server to a particular transaction, such that an assigned server processes the transaction to which it is assigned. After the processing of a particular assignment by a server is completed, the facility 100 produces a completed transaction 121. The server assigned to process the completed transaction may exit the facility as a departing server 122. For example, human agents may leave their workstations for a break or at the end of the workday, while automatic response units may be deactivated for service or replacement. If the server does not leave the facility, however, it becomes a new arriving server 112, which is reassigned by the facility to a new transaction.

Figure 2 is a table diagram identifying a number of sample transaction types. Each row of the table 200 corresponds to a different transaction type and application for processing transactions of that type, and indicates an application number 201 uniquely identifying an application, a priority level 202, a transaction type 203 of the application, a service level goal 204, and a service level goal type 205. Each service level goal 204 is in turn made up of a performance criterion 206 and a minimum frequency 207. For example, row 221 of table 200 indicates that incoming calls from gold card customers have a priority level of 2. Their service level goal is to assign transactions within 60 seconds of their arrival (their performance criterion) for 90% of the transactions (their minimum frequency). This is a transaction assignment service level goal. Row 241 differs somewhat from the other rows, in that it represents a default activity for servers that are not needed for any application. As such, it has a

service level goal type of "none." When a server becomes available, the facility preferably offers the server first to the priority level 1 application identified in row 211, then to the priority level 2 applications identified in rows 221 and 222, then to the priority level 3 applications identified in rows 231-233, then to the priority level 100 application identified in row 241, until the offered server is accepted.

Figure 3 is a conceptual diagram showing the facility's approach to assigning available servers to the priority levels identified in Figure 2. In Figure 3, when each arriving server 301 becomes available to the facility, the facility offers the arriving server to one priority level after another until one of the priority levels accepts the offered server. The facility begins by offering the server to the highest priority level, that is, the priority level having the smallest number. In the example, highest priority level is 1. Priority level 1 (311) determines whether the offered server is needed to meet the service level goal of any of its applications. If so, priority level 1 accepts the offered server, and assigns it to one of its applications (discussed in greater detail below), which uses it to process an arriving transaction. After such processing is complete, the application outputs the completed transaction, and releases the server. At this point, the server may join the departing servers 302 that are no longer available to process transactions, or may follow route 303 to reenter the facility as an arriving server.

If priority level 1 does not accept the offered server, the facility next offers the server to priority level 2 (312). Again, if priority level 2 accepts the server, then one of the applications in priority level 2 uses the server to process one of its arriving transactions, and then that application releases the server for departure or reassignment. If priority level 2 does not accept the offered server, then the facility offers the server to priority level 3 (313).

If, after offering the server to priority levels 1-3, the server has not been accepted, the facility preferably offers the application to priority level 100 (314). Priority level 100 has a minimum priority, and contains a single "background" application, which always has transactions available. This background application is preferably provided to accept any offered servers that are not accepted by applications

in higher priority levels. As the net result of this process, the server is accepted and used by exactly one of the four priority levels, then released either to depart or to be reassigned.

Figure 4 is a flow diagram showing the steps preferably performed by the facility in order to assign an available server to one of the priority levels. In steps 5 401-404, the facility loops through each defined priority level, from the highest priority level to the lowest. In terms of the applications identified in Figure 2, this involves looping through priority levels 1, 2, 3, and 100, in that order. Step 401 marks the beginning of this loop, and step 404 marks its end. In step 402, the facility offers the 10 available server to the current priority level (*i.e.*, level 1 in the first loop through). The details of step 402 are discussed in greater detail below in connection with Figure 6. In step 403, if, in response to the offer of step 402, the server is accepted by the current priority level, then these steps conclude. If in step 402 the server is not accepted, the facility continues in step 404. In step 404, if additional, lower priority levels remain, 15 then the facility loops back to step 401 to offer the server to the next lower priority level. If no lower priority levels remain, these steps conclude. If no default application is defined for the facility, the facility may loop through all of the priority levels without assigning the available server to any of the priority levels. In this case, the facility preferably repeats these steps at a later time.

20 Figure 5 is a conceptual diagram showing the facility's approach to assigning available servers to the sample applications identified in Figure 2 within a particular priority level. Figure 5 shows the operation of distribution logic for priority level 3 as identified in the table shown in Figure 2, containing applications 4, 5, and 6. After an available agent has been offered to priority levels 1 and 2 and declined by both, 25 the available server 501 is offered to priority level 3. In response, the priority level 3 server distribution logic (1) prioritizes the applications in the service level in terms of their need for available servers, then tests for server demand in the applications in that priority order. In prioritizing, the server distribution logic 500 uses indications of the current or expected future service level of each of the applications 511-513. The server 30 distribution logic then offers the available server, in turn, to the three applications in the

priority level 521-523 in this order. As discussed below, each application, when it is offered the available server, determines whether the server is needed to meet its service level goal, and if so, accepts the server. If the available server is not needed by the current application to meet its service level goal, the current application declines the available server and the priority level server distribution logic offers it to the next application in the priority order. If the available server is accepted by one of the applications, it is used to process a transaction, then released as a departing server 503. On the other hand, if the available server is not accepted by any of the applications, the server is declined by the priority level 502, and can be offered to the next lower priority level.

Figure 6 is a flow diagram showing the steps preferably performed by the facility in order to offer an available server to a particular priority level, as described above for step 402 of Figure 4. In step 601, the facility establishes an order for the applications having the current priority level. The facility preferably omits step 601 for priority levels having only one application. In terms of the applications identified in Figure 2, for priority level 3, the facility would establish an order for the applications identified in rows 231, 232, and 233. In cases in which the server is needed by more than one application in the current priority level, the order established in step 601 determines which of these applications is permitted to use the server. The facility therefore establishes the sequence in a manner that is fair to all of the applications of the current priority level. In one preferred embodiment, the facility uses a round-robin approach for establishing this order, beginning with an initial random order and, for each offer, rotating the order one position in a consistent direction. In this embodiment, each application of the priority level has equal opportunity to accept servers offered to the priority level. In another preferred embodiment, the facility assesses the success of each application at reaching its service level goal, and establishes a sequence that favors the applications that are the furthest from reaching their service level goals. In this embodiment, the facility exerts a balancing force on the applications of the priority level, enabling underperforming applications to receive extra servers that can be used to improve the performance of these applications.

In steps 602-605 of Figure 6, the facility loops through each application in the order established in step 601. Step 602 marks the beginning of this loop, and step 605 marks its end. In step 603, the facility offers the available server to the current application. The performance of step 603 is discussed in greater detail below in conjunction with Figure 7. In step 604, if the current application accepts the server in response to the offer of step 603, then these steps conclude. If the current application does not accept the server, the facility continues in step 605. In step 605, if additional applications remain in the established order, then the facility loops back to step 602 to offer the available server to the next application. If no additional applications remain, these steps conclude.

Figure 7 is a flow diagram showing the steps preferably performed by the facility in order to accept or decline an available server offered to a particular application, as described above for step 603 of Figure 6. In step 701, the facility determines whether the server is needed to satisfy the current application's service level goal. In terms of the applications identified in row 221 of Figure 2, for the incoming gold calls application, the facility determines whether the available server is needed in order to assign incoming gold calls within 60 seconds of their arrival for 90% of such calls. In step 701, if it is determined that the available server is needed to satisfy the current application's service level goal, then the facility continues in step 702. In step 702, the facility accepts the available server to process transactions within the current application. When the application accepts the offered server, if any transactions are waiting for processing the accepted server is assigned to process one of the waiting transactions. On the other hand, if no transactions are waiting when the application accepts the server, the accepted server is maintained in an idle/reserved server queue until a new transaction arrives. When a new transaction arrives, a server in the idle/reserved server queue is assigned to process the new transaction. After step 702, these steps conclude. If in step 701 the available server is not needed, the facility continues in step 703. In step 703, the facility declines to accept the server for the current application, and these steps conclude. In one preferred embodiment, the determination of step 701 is expanded to also include one or more tests (not shown) for

the suitability of the available server to the transaction type of the current application. In this embodiment, each individual server has certain qualification standards, which are used as the basis for determining whether a particular server is well-suited to a particular transaction type.

5 Figure 8 is a workflow diagram showing the activity performed by the facility with respect to a single application. Block 801 shows that a server is offered to the application, and the application either accepts or declines based on such parameters as the current service level ("SL"), the number of transactions waiting in a transaction queue 804 to be assigned (" N_{QUEUE} "), and the number of servers waiting to be assigned
10 to a transaction in an idle/reserved server queue 802 (" N_{IDLE} "). If the server is accepted, it is added to the idle/reserved server queue 802. Block 803 shows that arriving transactions are queued in a transaction queue 804. The rate at which transactions arrive is monitored by SL control logic 807. As shown by block 805, any time the idle/reserved server queue contains at least one server and the transaction queue
15 contains at least one transaction, a server is removed from the idle/reserved server queue 802, a transaction is removed from the transaction queue 804 and thus the removed server is assigned to process the removed transaction. When a server is assigned to a transaction in this manner, the SL control logic 807 monitors the amount of time the transaction spent in the transaction queue 804. Transactions being processed
20 by servers are shown in block 806. After a server completes the processing of a transaction, the facility preferably reassigns the server in accordance with the steps shown in Figure 6. Also, when processing of the transaction is completed, the SL control logic 807 monitors the amount of time the server took to process the transaction. The SL control logic 807 calculates statistics relating to the processing of transactions
25 in the application, such as N_{QUEUE} , N_{IDLE} , and SL, that are used in block 801 to determine whether to accept or decline an offered server.

The facility may be implemented on a wide variety of hardware configurations. Figure 9 is a high-level block diagram showing an illustrative computer system upon which the facility of the present invention may execute. As shown in
30 Figure 9, the computer system 900 comprises a central processing unit (CPU) 910,

input/output devices 920, and a computer memory 930. Among the input/output devices is a storage device 921, such as a hard disk drive; and a computer-readable media drive 922, which can be used to install software products, including the facility, which are provided on a computer-readable medium, such as a CD-ROM. The
5 input/output devices 930 further include a connection 923 to servers and transaction sources from which transactions are received. For example, the connection 923 may be used to connect the computer system to one or more computer telephony interfaces, world wide web servers, or world wide web browsers. This connection is usable by the facility to detect server availability, assign servers to applications and transactions, and
10 monitor the processing of transactions by servers.

The memory 930 contains the facility 931 for distributing servers, which includes a priority level server distribution subsystem 932 whose operation is discussed above in conjunction with Figure 4, and an application level server distribution subsystem 933 whose operation is discussed above in connection with Figure 6. In
15 addition to the facility 931, the memory 930 preferably also contains logic for each application, such as application 1 934 and application 2 936. Included for each application is a server acceptance subsystem for determining whether to accept servers offered to the application, such as server acceptance subsystem 935.

While the facility is preferably implemented on a computer system
20 configured as described above, those skilled in the art will recognize that it may also be implemented on computer systems having different configurations. For example, the facility may be implemented on a computer system having different components than described above. Further, the functionality of the computer system 900 may be divided among several different computers in a distributed system. For example, application 1
25 (934) and application 2 (936) may reside on separate, geographically-distributed hardware configurations interconnected via a communications network. The facility may also be implemented on special-purpose computer systems, such as those in a call center. The facility may further be implemented without software in special-purpose hardware, using, for example, application-specific integrated circuits.

While the present invention has been shown and described with reference to preferred embodiments, it will be understood by those skilled in the art that various changes or modifications in form and detail may be made without departing from the scope of the invention. For example, though described in conjunction with telephone call transactions, the facility may be used to allocate servers for processing
5 virtually any kind of transaction. Further, the facility may maintain additional state information that obviates certain steps of the allocation process described above. Additionally, the facility preferably facilitates the use of service level metrics and service level goal types other than those discussed above.

CLAIMS

I claim:

- 1 1. A method in a computer system for distributing transaction
2 servers to process transactions of a plurality of different types, the transaction types
3 being organized into a plurality of priority levels ordered from highest to lowest, each
4 priority level having one or more transaction types, each transaction type further having
5 a service level goal reflecting the level of service sought for transactions of the
6 transaction type, the method comprising the steps of:
7 identifying transaction servers when they become available to process
8 transactions;
9 for each identified transaction server, in response to the identification of
10 the transaction server:
11 for each priority level, from the highest priority level to the
12 lowest priority level, until the identified transaction server is assigned to process a
13 transaction of a selected type:
14 determining whether assigning the identified transaction
15 server to a transaction of any of the transaction types of the priority level is necessary to
16 meet the service level goal for that transaction type;
17 if it is determined that assigning the identified transaction
18 server to a transaction of any of the transaction types of the priority level is necessary to
19 meet the service level goal for that transaction type, assigning the identified transaction
20 server to process a transaction of one of the transaction types of the priority level; and
21 after the assigned transaction server has completed processing of
22 the transaction to which it was assigned, reidentifying the assigned transaction server as
23 being available to process additional transactions,
24 such that, for transaction types of a higher priority level relative to transaction types of a
25 lower priority level, an available transaction server cannot be assigned to process a
26 transaction of the transaction type of the lower priority level when the service level goal
27 for the transaction type of the higher priority level is not being satisfied, thus ensuring

28 that transaction servers are distributed in a manner that satisfies service level goals for
29 transaction types in order from the transaction types of the highest priority level toward
30 the transaction types of the lowest priority level.

1 2. The method of claim 1 wherein at least a portion of the
2 transaction types correspond to telephone calls, and wherein the step of identifying a
3 transaction server identifies transaction servers for conducting telephone calls.

1 3. The method of claim 2 wherein the identifying step identifies
2 human agents.

1 4. The method of claim 2 wherein the identifying step identifies
2 automatic telephone response units.

1 5. The method of claim 1 wherein at least a portion of the
2 transaction types correspond to web hits, and wherein the step of identifying a
3 transaction server identifies transaction servers for processing web hits.

1 6. The method of claim 5 wherein the identifying step identifies
2 human agents.

1 7. The method of claim 5 wherein the identifying steps identifies
2 automated web servers.

1 8. The method of claim 1 wherein the determining and assigning
2 steps comprise the steps of:
3 generating an order for the transaction types of the priority level, from
4 first to last;
5 for each transaction type of the priority level, from first to last, until the
6 identified transaction server is assigned to process a transaction of a selected type:
7 determining whether assigning the identified transaction server to
8 a transaction of the transaction type is necessary to meet the service level goal for the
9 transaction type; and

10 if it is determined that assigning the identified transaction server
11 to a transaction of the transaction type is necessary to meet the service level goal for the
12 transaction type, assigning the identified transaction server to process a transaction of
13 the transaction type.

1 9. The method of claim 8 wherein the generating step generates an
2 order reflecting the extent to which the service level goals of the transaction types of the
3 priority level presently satisfied.

1 10. The method of claim 8 wherein the generating step generates an
2 order for the transaction types different from the order generated by an immediately
3 preceding performance of the generating step.

1 11. A computer-readable medium whose contents cause a computer
2 system to distribute transaction servers to process transactions of a plurality of different
3 types, the transaction types being organized into a plurality of priority levels ordered
4 from highest to lowest, each priority level having one or more transaction types, each
5 transaction type further having a service level goal reflecting the level of service sought
6 for transactions of the transaction type, by performing the steps of:

7 identifying transaction servers when they become available to process
8 transactions;

9 for each identified transaction server, in response to the identification of
10 the transaction server:

11 for each priority level, from the highest priority level to the
12 lowest priority level, until the identified transaction server is assigned to process a
13 transaction of a selected type:

14 determining whether assigning the identified transaction
15 server to a transaction of any of the transaction types of the priority level is necessary to
16 meet the service level goal for that transaction type;

17 if it is determined that assigning the identified transaction
18 server to a transaction of any of the transaction types of the priority level is necessary to

19 meet the service level goal for that transaction type, assigning the identified transaction
20 server to process a transaction of one of the transaction types of the priority level; and
21 after the assigned transaction server has completed processing of
22 the transaction to which it was assigned, reidentifying the assigned transaction server as
23 being available to process transactions.

1 12. The computer-readable medium of claim 11 wherein contents of
2 the computer-readable medium further cause the computer system to perform the steps
3 of, as part of the determining and assigning steps:
4 generating an order for the transaction types of the priority level, from
5 first to last;
6 for each transaction type of the priority level, from first to last, until the
7 identified transaction server is assigned to process a transaction:
8 determining whether assigning the identified transaction server to
9 a transaction of the transaction type is necessary to meet the service level goal for the
10 transaction type; and
11 if it is determined that assigning the identified transaction server
12 to a transaction of the transaction type is necessary to meet the service level goal for the
13 transaction type, assigning the identified transaction server to process a transaction of
14 the transaction type.

1 13. An apparatus for distributing transaction servers to process
2 transactions of a plurality of different types, the transaction types being organized into a
3 plurality of priority levels ordered from highest to lowest, each priority level having one
4 or more transaction types, the apparatus comprising:
5 a service level goal attribution subsystem for attributing a service level
6 goal to each transaction type that reflects the level of service sought for transactions of
7 this type;
8 an identification subsystem that identifies transaction servers when they
9 become available to process transactions; and

10 a transaction server distribution subsystem that, in response to each
11 transaction server identification by the identification subsystem, assigns the identified
12 transaction server to process a transaction of a type that needs an additional transaction
13 server to meet its service level goal, and whose priority level is at least as high as that of
14 any other transaction type that needs an additional transaction server to meet its service
15 level goal,
16 the identification subsystem reidentifying assigned transaction servers when they
17 complete processing of the transactions to which they were assigned.

1 14. A computer memory containing a transaction server allocation
2 data structure usable to allocate an available transaction server to processing one of a
3 plurality of transaction types, the data structure comprising, for each transaction type:
4 an indication of the priority of transactions of the transaction type
5 relative to the priorities of transactions of other transaction types; and
6 an indication of a service level goal reflecting the level of service sought
7 for transactions of the transaction type,
8 such that the data structure may be used to allocate an available transaction server by,
9 for each transaction type, in descending order of the priorities indicated for the
10 transactions of the transaction types, determining whether the available transaction
11 server is needed in order to meet the service level goal indicated for the transaction
12 type, and assigning the available transaction server to the highest priority transaction
13 type for which the available transaction server is determined to be needed.

1 15. The computer memory of claim 14 wherein at least one pair of
2 priority indications contained in the data structure indicate the same level of priority.

1 16. The computer memory of claim 14 wherein at least a portion of
2 the service level goal indications indicate a service level goal specifying a target rate at
3 which transactions of the transaction type are to be assigned to available transaction
4 servers.

1 17. The computer memory of claim 14 wherein at least a portion of
2 the service level goal indications indicate a service level goal specifying a target
3 minimum amount of time in which transactions of the transaction type are to be
4 assigned to available transaction servers.

1 18. The computer memory of claim 14 wherein at least a portion of
2 the service level goal indications indicate a service level goal specifying a target rate at
3 which transactions of the transaction type are to be completed by the available
4 transaction servers.

1 19. The computer memory of claim 14 wherein at least a portion of
2 the service level goal indications indicate a service level goal specifying a target
3 maximum amount of time in which processing of transactions of the transaction type
4 are to be completed by the available transaction servers.

1 20. A method in a computer system for distributing available
2 transaction servers among a plurality of transaction processing applications each having
3 the same priority, each application further having a service level goal indicating the
4 service level sought for the application, the method comprising the steps of, when a
5 transaction server becomes available:

6 establishing an order for the applications, from highest to lowest;

7 for each application, in the established order, until the available
8 transaction server is assigned to an application;

9 offering the available transaction server to the application;

10 under the control of the application, determining whether the
11 available transaction server is needed by the application in order for the application to
12 meet its service level goal; and

13 if it is determined that the available transaction server is needed
14 by the application in order for the application to meet its service level goal, assigning
15 the available transaction server to the application,

16 such that each available transaction server is assigned to the application that is highest
17 in the established order that needs the available transaction server in order to meet its
18 service goal.

1 21. The method of claim 20, further comprising the step of, before
2 the establishing step, determining a current service level of each application that reflects
3 the performance of the application, and wherein the establishing step establishes an
4 order based upon a comparison of current service level to the service level goal for each
5 application.

1 22. The method of claim 20, further comprising the step of, before
2 establishing the step, determining an expected service level of each application that
3 reflects the projected future performance of the application, and wherein the
4 establishing step establishes an order based upon a comparison of the expected service
5 level to the service level goal for each application.

1 23. The method of claim 20 wherein each iteration of the establishing
2 step establishes an order that is rotated one position from the order established from the
3 previous iteration of the establishing step, such that the applications each have a
4 comparable opportunity to obtain available transaction servers.

1 24. The method of claim 20 wherein a selected one of the
2 applications has a server qualification standard, and wherein only a portion of the
3 transaction servers have qualification credentials satisfying the server qualification
4 standard, and wherein the determining step for the selected application includes the step
5 of determining whether the available transaction server has qualification credentials
6 satisfying the server qualification standard.

1 25. A computer-readable medium whose contents cause a computer
2 system to distribute available transaction servers among a plurality of transaction
3 processing applications each having the same priority, each application further having a

4 service level goal indicating the service level sought for the application, by performing
5 the steps of, when a transaction server becomes available:
6 establishing an order for the applications, from highest to lowest;
7 for each application, in the established order, until the available
8 transaction server is assigned to an application;
9 offering the available transaction server to the application;
10 under the control of the application, determining whether the
11 available transaction server is needed by the application in order for the application to
12 meet its service level goal; and
13 if it is determined that the available transaction server is needed
14 by the application in order for the application to meet its service level goal, assigning
15 the available transaction server to the application.

1 26. The computer-readable medium of claim 25 wherein the contents
2 of the computer readable medium further cause the computer system to perform the step
3 of, before the establishing step, determining a current service level of each application
4 that reflects the performance of the application, and wherein the establishing step
5 establishes an order based upon a comparison of current service level to the service
6 level goal for each application.

1 27. The computer readable medium of claim 25, further comprising
2 the step of, before establishing the step, determining an expected service level of each
3 application that reflects the projected future performance of the application, and
4 wherein the establishing step establishes an order based upon a comparison of the
5 expected service level to the service level goal for each application.

1 28. The computer-readable medium of claim 25 wherein each
2 iteration of the establishing step establishes an order that is rotated one position from
3 the order established from the previous iteration of the establishing step, such that the
4 applications each have a comparable opportunity to obtain available transaction servers.

1 29. The computer readable medium of claim 25 wherein a selected
2 one of the applications has a server qualification standard, and wherein only a portion of
3 the transaction servers have qualification credentials satisfying the server qualification
4 standard, and wherein the determining step for the selected application includes the step
5 of determining whether the available transaction server has qualification credentials
6 satisfying the server qualification standard.

1 30. A method in a computer system for maintaining a target service
2 level of a transaction processing application that uses reassignable transaction servers to
3 process transactions that arrive from a transaction source, the method comprising the
4 steps of:
5 receiving an offer of an additional transaction server;
6 determining the rate at which transactions are arriving from the
7 transaction source;
8 determining the number of transactions that have arrived from the
9 transaction source whose processing has not been assigned to a transaction server;
10 determining the rate at which processing of transactions is being
11 completed; and
12 based on the target service level, the determined rate at which
13 transactions are arriving from the transaction source, the determined number of
14 transactions that have arrived from the transaction source whose processing has not
15 been assigned to a transaction server, and the rate at which processing of transactions is
16 being completed, determining whether the offered additional transaction server is
17 needed to satisfy the target service level; and
18 accepting the offered additional transaction server to enable the offered
19 additional transaction server to be assigned to a transaction of the application only if it
20 is determined that the offered additional transaction server is needed to satisfy the target
21 service level.

1 31. The method of claim 30 wherein the application has an
2 idle/reserved server queue of additional transaction servers that have not yet been
3 accepted to process transactions, and wherein the accepting step includes the step of
4 adding the offered additional transaction server to the idle/reserved server queue.

1 32. The method of claim 30 wherein the application has a transaction
2 queue of transactions to whose processing no transaction server has yet been assigned,
3 and wherein the accepting step includes the step of assigning the offered additional
4 transaction server to process a particular transaction in the transaction queue.

1 33. The method of claim 30 wherein the step of determining whether
2 the offered additional transaction server is needed utilizes a quantitative model of
3 transaction processing within the application to determine whether the offered
4 additional transaction server is needed to satisfy the target service level.

1 34. The method of claim 30 wherein the step of determining whether
2 the offered transaction server is needed utilizes a mathematical formula modeling
3 transaction processing within the application to determine whether the offered
4 additional transaction server is needed to satisfy the target service level.

1 35. The method of claim 30 wherein the step of determining whether
2 the offered additional transaction server is needed utilizes a data structure indicating, for
3 a plurality of combinations of the rate at which transactions are arriving from the
4 transaction source, the determined number of transactions that have arrived from the
5 transaction source whose processing has not been assigned to a transaction server, and
6 the determined rate at which processing of transactions is being completed, whether an
7 offered additional transaction server is needed to satisfy the target service level.

1 36. A computer-readable medium whose contents cause a computer
2 system to maintain a target service level of a transaction processing application that
3 uses reassignable transaction servers to process transactions that arrive from a
4 transaction source by performing the steps of:

5 receiving an offer of an additional transaction server;

6 determining the rate at which transactions are arriving from the
7 transaction source;

8 determining the number of transactions that have arrived from the
9 transaction source whose processing has not been assigned to a transaction server;

10 determining the rate at which processing of transactions is being
11 completed; and

12 based on the target service level, the rate at which transactions are
13 arriving from the transaction source, the determined number of transactions that have
14 arrived from the transaction source whose processing has not been assigned to a
15 transaction server, and the determined rate at which processing of transactions is being
16 completed, determining whether the offered transaction server is needed to satisfy the
17 target service level; and

18 accepting the offer only if it is determined that the offered transaction
19 server is needed to satisfy the service level.

1 37. A method in a computer system for maintaining a target service
2 level of a transaction processing application that uses reassignable transaction servers to
3 process transactions that arrive from a transaction source, the method comprising the
4 steps of:

5 receiving an offer of an additional transaction server;

6 determining a current service level reflecting the current transaction
7 processing performance of the application;

8 comparing the determined current service level to the service level goal;
9 and

10 based on the result of the comparing step, determining whether the
11 offered transaction server is needed to satisfy the service level goal; and
12 accepting the offer only if it is determined that the offered transaction
13 server is needed to satisfy the service level goal.

1 38. The method of claim 37 wherein the application has an
2 idle/reserved server queue of transaction servers that have not yet been assigned to
3 process a transaction, and wherein the accepting step includes the step of adding the
4 offered transaction server to the idle/reserved server queue.

1 39. The method of claim 37 wherein the application has a transaction
2 queue of transactions to whose processing no transaction server has yet been assigned,
3 and wherein the accepting step includes the step of assigning the offered transaction
4 server to process a particular transaction in the transaction queue.

1 40. A method in a computer system for distributing reassignable
2 transaction processing servers among a plurality of transaction processing applications,
3 the method comprising the steps of:

4 defining for a first application of the plurality a transaction assignment
5 service level goal reflecting the performance sought for the first application in assigning
6 transactions to a server;

7 defining for a second application of the plurality a transaction
8 completion service level goal reflecting the performance sought for the second
9 application in completing the processing transactions;

10 when a server is available for distribution:

11 assessing the demand for an additional server by the first
12 application based on the transaction assignment service level goal defined for the first
13 application;

14 assessing the demand for an additional server by the second
15 application based on the transaction completion service level goal defined for the
16 second application;

17 comparing the assessments; and
18 based upon the results of the comparing step, selecting between
19 the first and second application for distribution of the available server.

1 41. The method of claim 40 wherein the first application is for
2 processing telephone call transactions, and wherein the step of defining a transaction
3 assignment level goal defines a transaction assignment service level goal reflecting the
4 performance sought in assigning telephone call transactions to servers.

1 42. The method of claim 41 wherein the second application is for
2 processing transactions other than telephone calls, and wherein the step of defining a
3 transaction completion service goal defines a transaction completion service goal
4 reflecting the performance sought in assigning transactions other than telephone calls to
5 servers.

1 43. A computer-readable medium whose contents cause a computer
2 system to distribute reassignable transaction processing servers among a plurality of
3 transaction processing applications by performing the steps of:

4 defining for a first application of the plurality a transaction assignment
5 service level goal;

6 defining for a second application of the plurality a transaction
7 completion service level goal;

8 when a server is available for distribution:

9 assessing the demand for an additional server by the first
10 application based on the transaction assignment service level goal defined for the first
11 application;

12 assessing the demand for an additional server by the second
13 application based on the transaction completion service level goal defined for the
14 second application;

15 comparing the assessments; and

- 16 based upon the results of the comparing step, selecting between the first
17 and second application for distribution of the available server.

Amendments to the claims have been filed as follows

1. A method in a computer system for distributing transaction servers
1 to process transactions of a plurality of different types, the transaction types being
2 organized into a plurality of priority levels ordered from highest to lowest, each priority
3 level having one or more transaction types, each transaction type further having a service
4 level goal reflecting the level of service sought for transactions of the transaction type,
5 the method comprising the steps of:
6 identifying transaction servers when they become available to process
7 transactions;
8 for each identified transaction server, in response to the identification of
9 the transaction server:
10 for each priority level, from the highest priority level to the lowest
11 priority level, until the identified transaction server is assigned to process a transaction of
12 a selected type:
13 determining whether assigning the identified transaction
14 server to a transaction of any of the transaction types of the priority level is necessary to
15 meet the service level goal for that transaction type;
16 if it is determined that assigning the identified transaction
17 server to a transaction of any of the transaction types of the priority level is necessary to
18 meet the service level goal for that transaction type, assigning the identified transaction
19 server to process a transaction of one of the transaction types of the priority level; and
20 after the assigned transaction server has completed processing of
21 the transaction to which it was assigned, reidentifying the assigned transaction server as
22 being available to process additional transactions,
23 such that, for transaction types of a higher priority level relative to transaction types of a
24 lower priority level, an available transaction server cannot be assigned to process a

26 transaction of the transaction type of the lower priority level when the service level goal
27 for the transaction type of the higher priority level is not being satisfied, thus ensuring
28 that transaction servers are distributed in a manner that satisfies service level goals for
29 transaction types in order from the transaction types of the highest priority level toward
30 the transaction types of the lowest priority level.

2. The method of claim 1 wherein the determining and assigning
1 steps comprise the steps of:
2 generating an order for the transaction types of the priority level, from first
3 to last;
4 for each transaction type of the priority level, from first to last, until the
5 identified transaction server is assigned to process a transaction of a selected type:
6 determining whether assigning the identified transaction server to a
7 transaction of the transaction type is necessary to meet the service level goal for the
8 transaction type; and
9 if it is determined that assigning the identified transaction server to
10 a transaction of the transaction type is necessary to meet the service level goal for the
11 transaction type, assigning the identified transaction server to process a transaction of the
12 transaction type.

3. The method of claim 2 wherein the generating step generates an
1 order reflecting the extent to which the service level goals of the transaction types of the
2 priority level presently satisfied.

4. The method of claim 2 wherein the generating step generates an
1 order for the transaction types different from the order generated by an immediately
2 preceding performance of the generating step.

5. An apparatus for distributing transaction servers to process
1 transactions of a plurality of different types, the transaction types being organized into a
2 plurality of priority levels ordered from highest to lowest, each priority level having one
3 or more transaction types, the apparatus comprising:

4 a service level goal attribution subsystem for attributing a service level
5 goal to each transaction type that reflects the level of service sought for transactions of
6 this type;

7 an identification subsystem that identifies transaction servers when they
8 become available to process transactions; and

9 a transaction server distribution subsystem that, in response to each
10 transaction server identification by the identification subsystem, assigns the identified
11 transaction server to process a transaction of a type that needs an additional transaction
12 server to meet its service level goal, and whose priority level is at least as high as that of
13 any other transaction type that needs an additional transaction server to meet its service
14 level goal,

15 the identification subsystem reidentifying assigned transaction servers when they
16 complete processing of the transactions to which they were assigned.

6. A computer memory containing a transaction server allocation data
1 structure usable to allocate an available transaction server to processing one of a plurality
2 of transaction types, the data structure comprising, for each transaction type:

3 an indication of the priority of transactions of the transaction type relative
4 to the priorities of transactions of other transaction types; and

5 an indication of a service level goal reflecting the level of service sought
6 for transactions of the transaction type,

7 such that the data structure may be used to allocate an available transaction server by, for
8 each transaction type, in descending order of the priorities indicated for the transactions

10 of the transaction types, determining whether the available transaction server is needed in
11 order to meet the service level goal indicated for the transaction type, and assigning the
12 available transaction server to the highest priority transaction type for which the available
13 transaction server is determined to be needed.

7. The computer memory of claim 6 wherein at least one pair of
1 priority indications contained in the data structure indicate the same level of priority.

8. The computer memory of claim 6 wherein at least a portion of the
1 service level goal indications indicate a service level goal specifying a target rate at which
2 transactions of the transaction type are to be assigned to available transaction servers.

9. The computer memory of claim 6 wherein at least a portion of the
1 service level goal indications indicate a service level goal specifying a target minimum
2 amount of time in which transactions of the transaction type are to be assigned to
3 available transaction servers.

10. The computer memory of claim 6 wherein at least a portion of the
1 service level goal indications indicate a service level goal specifying a target rate at which
2 transactions of the transaction type are to be completed by the available transaction
3 servers.

11. The computer memory of claim 6 wherein at least a portion of the
1 service level goal indications indicate a service level goal specifying a target maximum
2 amount of time in which processing of transactions of the transaction type are to be
3 completed by the available transaction servers.

12. A computer-readable medium whose contents cause a computer
1 system to distribute available transaction servers among a plurality of transaction
2 processing applications each having the same priority, each application further having a

4 service level goal indicating the service level sought for the application, by performing
5 the steps of, when a transaction server becomes available:
6 establishing an order for the applications, from highest to lowest;
7 for each application, in the established order, until the available transaction
8 server is assigned to an application;
9 offering the available transaction server to the application;
10 under the control of the application, determining whether the
11 available transaction server is needed by the application in order for the application to
12 meet its service level goal; and
13 if it is determined that the available transaction server is needed by
14 the application in order for the application to meet its service level goal, assigning the
15 available transaction server to the application.

13. The computer-readable medium of claim 12 wherein the contents
1 of the computer readable medium further cause the computer system to perform the step
2 of, before the establishing step, determining a current service level of each application
3 that reflects the performance of the application, and wherein the establishing step
4 establishes an order based upon a comparison of current service level to the service level
5 goal for each application.

14. The computer readable medium of claim 12, further comprising the
1 step of, before establishing the step, determining an expected service level of each
2 application that reflects the projected future performance of the application, and wherein
3 the establishing step establishes an order based upon a comparison of the expected
4 service level to the service level goal for each application.

15. The computer-readable medium of claim 12 wherein each iteration
1 of the establishing step establishes an order that is rotated one position from the order

3 established from the previous iteration of the establishing step, such that the applications
4 each have a comparable opportunity to obtain available transaction servers.

16. The computer readable medium of claim 12 wherein a selected one
1 of the applications has a server qualification standard, and wherein only a portion of the
2 transaction servers have qualification credentials satisfying the server qualification
3 standard, and wherein the determining step for the selected application includes the step
4 of determining whether the available transaction server has qualification credentials
5 satisfying the server qualification standard.

17. A method in a computer system for distributing reassignable
1 transaction processing servers among a plurality of transaction processing applications,
2 the method comprising the steps of:

3 defining for a first application of the plurality a transaction assignment
4 service level goal reflecting the performance sought for the first application in assigning
5 transactions to a server;

6 defining for a second application of the plurality a transaction completion
7 service level goal reflecting the performance sought for the second application in
8 completing the processing transactions;

9 when a server is available for distribution:

10 assessing the demand for an additional server by the first
11 application based on the transaction assignment service level goal defined for the first
12 application;

13 assessing the demand for an additional server by the second
14 application based on the transaction completion service level goal defined for the second
15 application;

16 comparing the assessments; and

17 based upon the results of the comparing step, selecting between the
18 first and second application for distribution of the available server.

1 18. The method of claim 17 wherein the first application is for
2 processing telephone call transactions, and wherein the step of defining a transaction
3 assignment level goal defines a transaction assignment service level goal reflecting the
4 performance sought in assigning telephone call transactions to servers.

1 19. The method of claim 18 wherein the second application is for
2 processing transactions other than telephone calls, and wherein the step of defining a
3 transaction completion service goal defines a transaction completion service goal
4 reflecting the performance sought in assigning transactions other than telephone calls to
5 servers.



4

Application No: GB 9905212.8
Claims searched: 1-19

Examiner: Ben Micklewright
Date of search: 27 September 1999

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.Q): G4A (AFN)

Int Cl (Ed.6): G06F (9/46 17/60)

Other: Online: WPI, EPODOC, JAPIO

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	US 5 504 894 (IBM) See whole document, e.g. column 2 line 20 to column 3 line 53	1-19
X	US 5 473 773 (IBM) See whole document, e.g. column 2 line 18 to column 3 line 61	1-17

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.