

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2005242593 B2**

- (54) Title
Queuing system, method and computer program product for managing the provision of services over a communications network
- (51) International Patent Classification(s)
H04L 29/06 (2006.01) **H04L 29/08** (2006.01)
G06Q 10/00 (2006.01)
- (21) Application No: **2005242593** (22) Date of Filing: **2005.05.12**
- (87) WIPO No: **WO05/112389**
- (30) Priority Data
- | (31) Number | (32) Date | (33) Country |
|------------------|-------------------|--------------|
| 0410829.6 | 2004.05.14 | GB |
| 0500801.6 | 2005.01.14 | GB |
- (43) Publication Date: **2005.11.24**
(44) Accepted Journal Date: **2009.08.06**
- (71) Applicant(s)
Orderly Mind Limited
- (72) Inventor(s)
King, Matt
- (74) Agent / Attorney
Madderns, Level 1 64 Hindmarsh Sq, Adelaide, SA, 5000
- (56) Related Art
US 2003/0163737 A1 (ROSKIND) 28 August 2003
US 2003/0023734 A1 (MARTIN et al) 30 January 2003
US 2003/0039350 A1 (HOLMEN) 27 February 2003

(WO/2005/112389) QUEUING SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR MANAGING THE PROVISION OF SERVICES OVER A COMMUNICATIONS NETWORK

Biblio. Data Description Claims National Phase Notices Documents

Latest bibliographic data on file with the International Bureau

Publication Number: WO/2005/112389 **International Application No.:** PCT/GB2005/001854
Publication Date: 24.11.2005 **International Filing Date:** 12.05.2005

Int. Class.: *H04L 29/06* (2006.01)

Applicant: KING, Matt [GB/GB]; 36 Thornton Hill, Exeter, Devon EX4 3BH (GB).


Inventor: **KING, Matt** [GB/GB]; 36 Thornton Hill, Exeter, Devon EX4 3BH (GB).

Agent: GILL JENNINGS & EVERY LLP; Broadgate House, 7 Eldon Street, London EC2M 7LH (GB).

Priority Data: 0410829.6 14.05.2004 GB
0500801.6 14.01.2005 GB

Title: QUEUING SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR MANAGING THE PROVISION OF SERVICES OVER A COMMUNICATIONS NETWORK

Abstract: The present invention provides a method for managing requests for service over a communications network. The method of the invention comprises the steps of: receiving a request for service from a customer terminal at a queue server via the communications network; allocating a queue identifier to the request for service; sending the queue identifier to the customer terminal; receiving the queue identifier from the customer terminal at the queue server as part of a subsequent request for service; performing a comparison between the queue identifier and queue status information; and forwarding the request for service to the service host in accordance with the result of the comparison. Preferably, if sufficient resources are available, the request for service is forwarded directly to the service host without entering a queue. If there are insufficient resources, the request for service is held in an automatically managed queue and the risk of catastrophic failure is eliminated.



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

African Regional Intellectual Property Org. (ARIPO) (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, T
UG, ZM, ZW)

Eurasian Patent Organization (EAPO) (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM)

European Patent Office (EPO) (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR)

African Intellectual Property Organization (OAPI) (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Publication Language: English (EN)

Filing Language: English (EN)



**QUEUEING SYSTEM, METHOD AND COMPUTER PROGRAM
PRODUCT FOR MANAGING THE PROVISION OF SERVICES
OVER A COMMUNICATIONS NETWORK**

5 **Field of the Invention**

The present invention relates to the provision of services over a communications network having limited bandwidth, due to a limitation of the network or limitation on server resources.

10 **Background to the Invention**

Any service provided over a communications network will have limited bandwidth, resulting in a maximum number of customers that can be served per minute. The bandwidth may be limited for technical reasons, such as web service speed or the number of incoming phone lines, or may be limited because there are
15 simply not enough operators to handle the demand for service.

When demand is less than the maximum possible service rate, customers can be served immediately. However, when demand is greater than the maximum possible service rate, problems can occur unless a traffic management system is employed. Excessive demand often occurs in e-commerce when there is a very high interest in
20 a particular product, which may be available in limited quantities when it first goes on sale. A typical example is the selling of concert tickets using an e-commerce system. Fans, knowing that tickets are limited, will all try to use the system as soon as the tickets go on sale, creating a demand "spike" that may well be above the maximum transaction rate that the system can cope with.

25 The present invention aims to provide a system and method which deals with excessive demands of this type in a way which is acceptable to customers whilst being economically viable.

An e-commerce transaction is a process by which a customer obtains access to a product or service through the processing of payment information, usually
30 consisting of credit card details. A customer participating in an on-line purchase typically follows the following steps:

1. The customer browses a website, making requests and reading responses using web browsing software. The pages of the website are normally transferred from the web server to the customer's client machine using unencrypted hypertext transfer
35 protocol (HTTP).

2. When they have made a choice, the customer requests a payment page, usually by pressing a "buy" button, containing form fields in which they can enter their credit/debit card information. This payment page is normally transferred over the network using secure HTTP (HTTPS).

5 3. The customer fills in his/her details, and these are sent using HTTPS to a payment gateway. Frequently, the pending transaction is stored in a database. The payment gateway forwards the transaction details onto a dedicated credit card network, such as VISA [RTM], for processing. Assuming that the transaction is authorised, the transaction is stored in the payment gateway's database. The transaction may also be
10 stored on databases attached to the payment page and the original web server.

4. The customer is presented with a response indicating that his or her transaction is successful. Usually a confirmation email is also sent to the customer.

At a later time, the goods will be sent to the customer and the amount deducted from their credit card or bank account. This is known as the confirmation of the
15 authorisation attained in step 3, and is usually achieved using a manual process.

In some configurations, the website, payment page and payment gateway all reside on different machines. In other configurations, the website and payment page may reside on the same machine with the payment gateway on a different machine. Alternatively, the payment page and payment gateway may reside on the same
20 machine.

The number of users who can successfully complete this path concurrently is limited by the bandwidth of the system, which in turn is limited by the bandwidth of the most resource intensive step. This is usually step 3, as this contains multiple database writes and always contains at least one connection to an external credit card network.

25 Typically the payment gateway will be able to cope with every requested transaction successfully up to an optimum rate of transactions per minute. As requested rates increase beyond this, some transactions will fail, resulting in frustrated users. However, the number of successful transactions per minute will continue to increase up to a maximum rate. Beyond this maximum request rate the number of
30 successful transactions falls, as the server has to spend ever increasing resources denying new incoming requests. This behaviour is illustrated in Figure 1. The optimal rate is indicated at 100 successful transactions per minute. Beyond this point the number of successful transactions per minute is less than the number of requested transactions per minute. However, the number of successful transactions per minute
35 reaches a maximum at about 120 successful transactions per minute.

Frequently, maximum transaction volumes are expressed in terms of number of simultaneous transactions. This can be calculated with the following formula from the maximum transaction rate as follows, provided the average duration of each transaction is also known:

$$5 \quad R = S * (60/T)$$

Where R is the maximum number of transactions per minute, T is the duration of each transaction in seconds (typically between 6 and 8 seconds), and S is the maximum number of simultaneous transactions. Note that this value for R is an upper limit, as the calculation assumes that each new transaction occurs immediately after a previous one completes. In practice this is unlikely, and "optimum" transaction rates are typically less than one third of this value.

Payment gateways may have an optimum transaction rate of about 100 successful transactions per minute, although most payment gateways are much less efficient than this. The curve shown in Figure 1 also applies to any web server, though the value of the optimal rate will depend on the number of database writes, the complexity and length of the page served and the transport protocol that is used. The HTTPs protocol, being encrypted, takes much more processing power for the same rate of page requests/response pairs than ordinary HTTP for example.

As the requested transaction rate at the payment gateway increases beyond the optimum level, users of the system will see their requests time out or be otherwise denied. A user intent on buying the product for sale will then cause the request to be re-submitted, by clicking on the submit or buy button in the form again or by hitting reload. This causes an additional request to be sent to the payment gateway, which must also be dealt with, further increasing the server load. For example, referring to Figure 1, if 200 users make requests in one minute, around 90 of those will experience failed transactions. In the next minute, assuming all 90 of these users repeat their request, and another 200 new users enter the system, there will be 290 requests to the server. Only around 50 of these will experience successful transactions, giving 440 requests for the next minute and so on until catastrophic failure results.

30 Furthermore, as the payment gateway fails, users may start to jam up the earlier stages of the e-commerce path, resulting in saturated payment pages (stage 2) and eventually a saturated web server (stage 1). This problem is particularly significant if all these stages reside on the same physical machine.

As stated above, this situation is particularly serious when there is a very high
35 interest in the particular product which is available in limited quantities and goes on
sale at a particular time. A “spike” of customer requests is formed which is well above

the maximum transaction rate the system can cope with. Catastrophic failure then propagates backward through the system. Catastrophic failure will last as long as customers continue to submit requests and for popular events can last for many hours.

This creates a number of problems. Firstly, the experience of buying a product or service is protracted and is extremely frustrating, creating negative publicity for the service providers. Customers become angry that their time is being wasted. Secondly, customers perceive that the assignment of successful transactions is random or based on the whims of the networking hardware, leading to intense dissatisfaction. Thirdly, the server load means that attempts to modify the pages to improve the system “on the fly” may not succeed. If the high load has been anticipated, certain checks (such as credit card authorisation) may be postponed until after the 'sale' to increase performance, causing further work and uncertainty as to the number of tickets actually sold. The server load means that attempts must be made to increase capacity at an additional cost to the vendor and ultimately to the customer. Additional staff must also be hired to cope with the influx of orders and these staff members will have less time to deal with dissatisfied customers.

A simple but expensive solution to some of these problems is to increase the bandwidth of the most limited part of this system, usually the payment gateway, by adding additional servers in parallel. This is called scaling up. In practice, if the payment gateway is the bottleneck for the process, it is in general not possible to scale up its capacity to meet arbitrary demand rates. This is for two reasons. Firstly, the bandwidth of the credit card network is usually limited and secondly, the bandwidth of the database system that stores the transactions is limited. Database servers are very expensive to scale and there are often technical reasons why the number of database servers cannot be increased beyond two.

Whilst adding servers may allow a vendor in practice to double the transaction rate, the actual level of demand for a particular product cannot be accurately predicted and the risk of catastrophic failure for very high interest sales is not reduced. All that is achieved is further expense both in hardware and in staffing costs.

Finally, in real e-commerce, servers more frequently fail for reasons other than excess demand. Examples include network or power outages, hardware faults, system updates or reboots, and software errors. In these cases too, customers are prevented from obtaining offered products and services, frequently being presented with an error message instead of the desired successful transaction. Unless a service provider is alerted to a failure and performs the necessary repairs, these customers may be lost.

The present invention aims to reduce or eliminate some or all of the above problems, in a way that is satisfactory both to customers and is economically sensible for the vendors.

Summary of the Invention

Aspect 1. A method for managing requests for service to a service host over a communications network, the method comprising the steps of:

receiving from a customer terminal via the communications network a request for service to be supplied by the service host;

allocating a queue identifier to the request for service;

receiving from the customer terminal a subsequent request for service to be supplied by the service host;

retrieving the queue identifier for the request for service;

performing a comparison between the queue identifier and queue status information; and

passing to the service host the request for service in accordance with the result of the comparison,

wherein the steps of the method are performed at a queue server and the step of allocating the queue identifier to the request for service is performed independently of whether the request for service is currently eligible for service by the service host.

Aspect 2. A method according to aspect 1, further comprising the step of performing an initial comparison between the queue identifier and queue status information after the queue identifier has been allocated to the request for service.

Aspect 3. A method according to aspect 2, further comprising the step of notifying the customer terminal that the request has been enqueued after performing the initial comparison.

Aspect 4. A method according to aspect 3, further comprising the step of providing the customer terminal with access to the queue identifier.

Aspect 5. A method according to any preceding aspect, wherein the queue identifier is a number and is allocated on a sequential basis.

Aspect 6. A method according to any one of aspects 1 to 5, wherein the communications network is a telephone network.

Aspect 7. A method according to aspect 6, wherein the queue identifier is stored in a database and is associated with a telephone number of the customer terminal.

Aspect 8. A method according to aspect 7, wherein the queue identifier is retrieved from the database following the subsequent request for service.

Aspect 9. A method according to any one of aspects 6 to 8, wherein the step of passing the request for service to the service host comprises connecting a call from the customer terminal to the service host.

Aspect 10. A method according to any one of aspects 1 to 5, wherein the communications network is the Internet.

Aspect 11. A method according to aspect 10, wherein the queue identifier is retrieved from the subsequent request for service.

Aspect 12. A method according to aspect 10 or aspect 11, wherein the step of passing the request for service to the service host comprises providing the customer terminal access to a target page on the service host.

Aspect 13. A method according to aspect 4 or any one of aspects 5 to 12 when dependent on aspect 4, further comprising the step of providing the customer terminal with access to additional information relating to the queue.

Aspect 14. A method according to aspect 13, wherein the additional information includes a service number, the service number being an indication of the queue identifier currently being served.

Aspect 15. A method according to aspect 14, wherein the service number is automatically incremented at a constant rate.

Aspect 16. A method according to aspect 14, wherein the service number is automatically incremented, the rate of increment being varied in dependence on the number of requests for service passed to the service host.

Aspect 17. A method according to aspect 14, wherein the service number is incremented in response to messages from the service host.

Aspect 18. A method according to aspect 4 or any one of aspects 5 to 17 when dependent on aspect 4, wherein access to information specific to the customer terminal is provided with access to the queue identifier.

Aspect 19. A method according to any preceding aspect, further including the step of checking at the service host that the request for service has been passed from the queue server.

Aspect 20. A method according to any preceding aspect, further including the step of automatically deleting the queue identifier and any encrypted string from the customer terminal following the provision of service by the service host.

Aspect 21. A method according to any preceding aspect, wherein requests for service are received by a plurality of queue servers, further comprising the step of allocating the request for service from the customer terminal to a particular queue server in accordance with a predetermined metric.

Aspect 22. A method as aspected in any preceding aspect, wherein the queue identifier corresponds to one of a plurality of queues.

Aspect 23. A method according to aspect 22, further comprising balancing customer load between the queues.

Aspect 24. A method according to aspect 22 or aspect 23, wherein each queue corresponds to a given domain within the communications network.

Aspect 25. A method as aspected in aspect 24, wherein each domain corresponds to a geographical area.

Aspect 26. A method as aspected in aspect 24 or aspect 25, further comprising the steps of:
querying the customer terminal for a domain identifier;
receiving the domain identifier from the customer terminal; and
after the request for service is passed to the service host, checking that the domain identifier corresponds to domain information subsequently provided by the customer terminal when the request of service is passed to the service host.

Aspect 27. A method as aspected in any of the preceding aspects, further comprising the step of:
receiving a further request for service from a further customer terminal via the communications network, the further customer terminal being of a different type to the customer terminal.

Aspect 28. A method according to any preceding aspect, wherein the customer terminal is operated by a user, the method further comprising the steps of:

associating a customer identifier with each user;
receiving the customer identifier from the user; and
validating the customer identifier, thereby confirming that the user is a returning customer.

Aspect 29. A method as aspected in aspect 28, wherein the customer identifier is a customer terminal identifier that identifies a given customer terminal, and wherein the step of validating the customer identifier includes checking whether a further customer terminal, from which a request for service has been received, matches the previously allocated customer terminal identifier.

Aspect 30. A method as aspected in aspect 28 or aspect 29, wherein the step of associating the customer identifier with each user includes maintaining a queue database of customer identifiers, and wherein the step of validating the customer identifier includes comparing the customer terminal identifier for the customer terminal with each entry in the queue database and, where the customer terminal matches an entry in the queue database, confirming that the user is a returning customer.

Aspect 31. A method as aspected in aspect 28 or aspect 29, wherein the step of associating the customer identifier with each user includes obtaining further customer identification information, and wherein the step of validating the customer identifier includes validating the customer identifier against the further customer information.

Aspect 32. A method according to any one of aspects 28 to 31, wherein the step of associating the customer identifier with the user includes:

receiving customer information from the user;
generating a unique confirmation code for the user using the customer information; and,
allocating the confirmation code to the user.

Aspect 33. A method according to aspect 32, wherein the customer identifier received from the customer terminal includes the confirmation code and wherein the step of validating the customer identifier includes checking the confirmation code is a valid confirmation code, thereby allowing a plurality of customers to use the same customer terminal.

Aspect 34. A method according to any of the preceding aspects, further comprising the step of notifying the customer when the result of the comparison between the queue identifier and the queue status information indicates that the request for service should be passed to the service host.

Aspect 35. A method according to any one of the preceding aspects, wherein the queue status information indicates whether any or all subsystems of the service host have failed.

Aspect 36. A method according to any one of the preceding aspects, further comprising the steps of:

in advance of the comparison step, altering the queue status information to indicate that the requested service has been exhausted so that the request for service is prevented from being passed to the service host; and

when the requested service becomes available, altering the queue status information to indicate whether the request for service should be passed to the service host.

Aspect 37. A computer program product having computer executable code stored thereon for causing a computer to perform the method of any preceding aspect.

Aspect 38. A system for managing requests for service from a customer terminal to a service host over a communications network, the system comprising a server for receiving from the customer terminal via the communications network a request for service to be supplied by the service host, the server including:

means for receiving a request for service from the customer terminal;

means to allocate a queue identifier to the request for service;

means for generating queue status information;

means for retrieving the queue identifier for the request for service following a subsequent request for service from the customer terminal;

means to perform a comparison between the queue status information and the queue identifier; and

means for passing the request for service to the service host in dependence on the result of the comparison,

wherein the server is a queue server and the means for allocating the queue identifier to the request for service is adapted to allocate the queue identifier independently of whether the request for service is currently eligible for service by the service host.

Aspect 39. A system according to aspect 38, wherein the means for performing a comparison is adapted, in use, to perform an initial comparison between the queue identifier and queue status information after the means for allocating has allocated the queue identifier to the request for service.

Aspect 40. A system according to aspect 39, further comprising means for notifying the customer terminal that the request has been enqueued after performing the initial comparison.

Aspect 41. A system according to any one of aspects 38 to 40, further comprising means for providing the customer terminal with access to the queue identifier.

Aspect 42. A system according to any one of aspects 38 to 41, wherein queue identifiers are numbers and are allocated on a sequential basis.

Aspect 43. A system according to aspect 41 or aspect 42, wherein a subsequent request for service from the customer terminal is initiated automatically by code sent to the customer terminal when access to the queue identifier is provided.

Aspect 44. A system according to any one of aspects 38 to 43, wherein the communications network is a telephone network.

Aspect 45. A system according to aspect 44, further comprising means for storing the queue identifier in a database on behalf of the customer terminal and wherein the means for retrieving the queue identifier retrieves it from the database.

Aspect 46. A system according to aspect 44 or aspect 45, wherein the means for passing the request for service is adapted to connect a call from the customer terminal to the service host.

Aspect 47. A system according to any one of aspects 38 to 43, wherein the communications network is the Internet.

Aspect 48. A system according to aspect 47, wherein the means for retrieving the queue identifier is adapted to retrieve it from a subsequent request for service received from the customer terminal by the receiving means.

Aspect 49. A system according to aspect 47 or aspect 48, wherein the means for passing the request for service is adapted to provide the customer terminal with access to a target page on the service host.

Aspect 50. A system according to aspect 41 or any one of aspects 42 to 49 when dependent on aspect 41, wherein the means to provide access to the queue identifier to the customer terminal is adapted to provide access to additional information relating to the queue to the customer terminal with the queue identifier.

Aspect 51. A system according to aspect 50, wherein the additional information includes a service number, the service number being an indication of the queue identifier currently being served.

Aspect 52. A system according to aspect 51, wherein, in use, the queue server increments the service number automatically at a constant rate.

Aspect 53. A system according to aspect 51, wherein, in use, the queue server increments the service number automatically, the rate of increment being varied in dependence on the number of requests for service forwarded to the service host.

Aspect 54. A system according to aspect 41 or any one of aspects 42 to 53 when dependent on aspect 41, wherein the means to provide access to the queue identifier to the customer terminal is adapted to provide access to information specific to the customer terminal with the queue identifier.

Aspect 55. A system according to any one of aspects 38 to 54, further including means for checking at the service host that the request for service has been forwarded from the queue server.

Aspect 56. A system according to any one of aspects 38 to 55, further including means for automatically deleting the queue identifier and any encrypted string from the customer terminal following the provision of service by the service host.

Aspect 57. A system according to any one of aspects 38 to 56, further comprising means for automatically refusing requests for service after a predetermined number of requests for service have been forwarded to the service host.

Aspect 58. A system according to any one of aspects 38 to 57, comprising a plurality of queue servers, and means for allocating requests for service from customer terminals to a particular queue server in accordance with a predetermined metric.

According to a first aspect of the present invention, a method for managing requests for service to a service host over a communications network comprises the steps of:

receiving from a customer terminal via the communications network a request for service to be supplied by the service host;

allocating a queue identifier to the request for service;

receiving from the customer terminal a subsequent request for service to be supplied by the service host;

retrieving the queue identifier for the request for service;

performing a comparison between the queue identifier and queue status information; and

passing to the service host the request for service in accordance with the result of the comparison,

wherein the steps of the method are performed at a queue server and the step of allocating the queue identifier to the request for service is performed independently of whether the request for service is currently eligible for service by the service host.


Preferably, after the request for service is allocated a queue identifier an initial comparison between the queue identifier and queue status information is performed, and the request for service forwarded to the service host in accordance with the result of the initial comparison. If there are insufficient resources, the request for service is held in an automatically managed queue and the risk of catastrophic failure is eliminated. Preferably, queue identifiers are numbers and are allocated on a sequential basis.

The customer terminal may be notified that the request has been enqueued after performing the initial comparison and be provided with access to the queue identifier.

The subsequent request for service from the customer terminal may be initiated by the customer or may be initiated automatically. The method of the present invention allows customers to make a request for service and then disconnect. Their place in the queue is saved and when they reconnect they return to their place in the queue, possibly being served immediately.

Preferably, the communications network is the Internet. Preferably, access is provided by sending the queue identifier to the customer as a persistent cookie.

Further alternatives are equally preferable, they include providing access by sending the queue identifier to the customer terminal as a query string parameter; by sending the queue identifier within an HTML form; or by sending the queue identifier as a variable within a markup script language.



Alternatively, the communications network may be a telephone network. In which case, it is preferred that access to the queue identifier is provided by storing the queue identifier in a database on behalf of the customer terminal.

Preferably, additional information relating to the queue is sent to the customer terminal with the queue identifier. Preferably, the additional information includes a service number, the service number being an indication of the queue identifier currently being served. Preferably, the service number is automatically incremented at a constant rate. Alternatively, the service number may be automatically incremented, the rate of increment being varied in dependence on the number of requests for service forwarded to the service host. The rate at which the service number is incremented may be controlled by a system manager. In which case, the service number may be incremented in response to messages from the service host.

Preferably, at least some of the additional information is displayed on the client terminal.

Preferably, information specific to the customer terminal is sent with the queue identifier.

Preferably, the method further comprises the steps of:

encrypting the queue identifier, the information specific to the customer terminal and a secret phrase to form a first encrypted string;

sending the first encrypted string with the queue identifier to the customer terminal;

constructing a second encrypted string from the queue identifier and information specific to the customer terminal received as part of a subsequent request for service and the secret phrase; and

validating the subsequent request for service by comparing the first encrypted string with the second encrypted string. This prevents users jumping the queue.

Preferably, the method includes the step of checking at the service host that the request for service has been forwarded from the queue server. Preferably, the step of checking that the request for service has been forwarded from the queue server includes:

encrypting information specific to the request for service, information specific to the queue server and a secret phrase to form an encrypted queue server string;

sending the encrypted queue server string with the request for service to the service host; and

validating the request for service by comparing the encrypted queue server string with a service host encrypted string constructed at the service host using the

secret phrase and information specific to the request for service and information specific to the queue server. This prevents customers bypassing the queue server.

Preferably, the method further includes the step of automatically deleting the queue identifier and any encrypted string from the customer terminal following the provision of service by the service host. This prevents users reusing request identifiers.

Preferably, the method comprises the step of automatically terminating the method after a predetermined number of requests for service have been forwarded to the service host.

It may be that requests for service can be received by a plurality of queue servers. In this case the method includes the step of allocating the request for service from the customer terminal to a particular queue server in accordance with a predetermined metric. For example, the metric may be based on the geographical location of the customer terminal or in the case where more than one service can be provided, the particular service requested.

The method may further comprise the steps of, in advance of the receipt of the request for service receiving a group registration request from a first member; allocating a unique group identifier to the group registration request; sending the group identifier to the first member, thereby adding the first member to a registered group; receiving a subsequent group registration request from a further member, the subsequent group registration request including the group identifier allocated to the group registration request; and adding the further member to the registered group.

Where the request for service from the customer terminal includes the group identifier, thereby identifying the customer terminal as a member of the registered group, the method may further comprise, in advance of allocating the queue identifier, checking whether there is a queue identifier associated with the group identifier.

Advantageously, the step of allocating a queue identifier may include, where no queue identifier has yet been associated with the group identifier, associating a new queue identifier with the group identifier and sending the newly associated queue identifier to the customer terminal. Furthermore, where a group identifier is found to have an associated queue identifier, the step of allocating a queue identifier may include retrieving the associated queue identifier and sending the retrieved associated queue identifier to the customer terminal.

Where the subsequent request for service includes both the queue identifier and the group identifier, thereby identifying the customer terminal as a member of the registered group, the method may preferably further comprise:

receiving a further initial request for service from a further customer terminal at the queue server via the communications network, the further initial request including the group identifier, thereby identifying the further customer terminal as a member of the registered group;

- 5 allocating the queue identifier associated with the request for service from the customer terminal to the further request for service from the further customer terminal;
sending the queue identifier to the further customer terminal;
receiving the queue identifier from the further customer terminal at the queue server as part of a further subsequent request for service;
10 performing a comparison between the queue identifier and queue status information; and
forwarding the further request for service to the service host in accordance with the result of the comparison.

- Each member of a registered group can then be assured of the opportunity to
15 access the requested service substantially simultaneously with the group member/requester closest to the "front" of the queue.

In a preferred embodiment, the queue identifier may correspond to one of a plurality of queues. The method may then further comprise balancing customer load between the queues.

- 20 Each such queue may correspond to a given domain within the communications network. The given domain may correspond to a geographical area. Alternatively, each domain may correspond to a given demographic component.

- The method may further comprise: querying the customer terminal for a domain identifier; receiving the domain identifier from the customer terminal; and after the
25 request for service is forwarded to the service host, checking that the domain identifier corresponds to domain information subsequently provided by the customer terminal when the request of service is forwarded to the service host.

- It is preferred that the method further includes receiving a further request for service from a further customer terminal via the communications network, the further
30 customer terminal being of a different type to the customer terminal.

The customer terminal is operated by a user. Advantageously, the method may further comprise associating a customer identifier with each user; receiving the customer identifier from the user; and validating the customer identifier, thereby confirming that the user is a returning customer.

- 35 Preferably, the customer identifier is a customer terminal identifier that identifies a given customer terminal, and the step of validating the customer identifier includes

checking whether a further customer terminal, from which a request for service has been received, matches the previously allocated customer terminal identifier.

Furthermore, the step of associating the customer identifier with each user may include maintaining a queue database of customer identifiers, and the step of validating
5 the customer identifier may include comparing the customer terminal identifier for the customer terminal with each entry in the queue database and, where the customer terminal matches an entry in the queue database, confirming that the user is a returning customer.

Alternatively, the step of associating the customer identifier with each user may
10 include obtaining further customer identification information, and the step of validating the customer identifier may include validating the customer identifier against the further customer information.

The step of associating the customer identifier with the user may also include:
receiving customer information from the user; using the customer information,
15 generating a unique confirmation code for the user; and allocating the confirmation code to the user. In the case where the customer identifier received from the customer terminal includes the confirmation code, the step of validating the customer identifier preferably includes checking the confirmation code is a valid confirmation code, thereby allowing a plurality of customers to use the same customer terminal.

It is preferred that the method further comprises querying the customer terminal
20 for personally identifiable information; receiving the personally identifiable information from the customer terminal; and, after the request for service is forwarded to the service host, checking that the personally identifiable information corresponds to personally identifiable information provided by the customer terminal when the request
25 for service is forwarded to the service host.

The method may further comprise the step of notifying the customer when the result of the comparison between the queue identifier and the queue status information indicates that the request for service should be forwarded.

The queue status information preferably indicates whether any or all
30 subsystems of the service host have failed.

The method may further comprise the steps of, in advance of the comparison step, altering the queue status information to indicate that the requested service has been exhausted so that the request for service is prevented from being forwarded to the service host; and, when the requested service becomes available, altering the
35 queue status information to indicate whether the request for service should be forwarded to the service host.

According to a second aspect of the present invention, a system for managing requests for service from a customer terminal to a service host over a communications network comprises a queue server for receiving a request for service from the customer terminal via the communications network, the queue server including:

5 means for receiving a request for service from the customer terminal;
means to allocate a queue identifier to the request for service;
means for generating queue status information;
means for retrieving the queue identifier for the request for service following a subsequent request for service from the customer terminal;
0 means to perform a comparison between the queue status information and the queue identifier;
and
means for passing the request for service to the service host in dependence on the result of the comparison,

wherein the means for allocating the queue identifier to the request for service is adapted to
5 allocate the queue identifier independently of whether the request for service is currently eligible for service by the service host.

Preferably, in use, after the request for service is allocated a queue identifier the means to perform a comparison performs an initial comparison between the queue identifier and queue status information, and the means for forwarding the request for service forwards the request for service to
10 the service host in accordance with the result of the initial comparison. If there are insufficient resources, the system of the present invention holds the request for service in an automatically managed queue and the risk of catastrophic failure is eliminated. Preferably, queue identifiers are numbers and are allocated on a sequential basis.


The system may further comprise means for notifying the customer terminal that the request
25 has been enqueued after performing the initial comparison and means for providing the customer terminal with access to the queue identifier.

The subsequent request for service from the customer may be initiated automatically by code sent to the customer terminal with the queue identifier. The system of the present invention allows customers to make a request for service and then disconnect. Their place in the queue is saved and
30 when they reconnect they return to their place in the queue.

Preferably, the communications network is the Internet. Preferably, the means for providing the customer terminal with access to the queue identifier sends the queue identifier to the customer terminal as a persistent cookie.

Equally preferably, the means for providing the customer terminal with access to the queue
35 identifier may send the queue identifier to the customer terminal as a query string parameter. Alternatively, it may send the queue identifier to the customer terminal within an HTML form or as a variable within a markup script language.

The communications network may alternatively be a telephone network. Preferably, the means for providing the customer terminal with access to the queue identifier stores the queue identifier in a database on behalf of the customer terminal.



Preferably, the means to send the queue identifier to the customer terminal is adapted to send additional information relating to the queue to the customer terminal with the queue identifier. Preferably, the additional information includes a service number, the service number being an indication of the queue identifier currently being served. Preferably, in use, the queue server increments the service number automatically at a constant rate. Preferably, in use, the queue server increments the service number automatically, the rate of increment being varied in dependence on the number of requests for service forwarded to the service host. Ideally, the rate at which the service number is incremented can be controlled.

10 Preferably, the means to send the queue identifier to the customer terminal is adapted to send information specific to the customer terminal with the queue identifier.

Preferably, the system further comprises:

15 means for encrypting the queue identifier, the information specific to the customer terminal and a secret phrase to form a first encrypted string;

means for sending the first encrypted string with the queue identifier to the customer terminal;

20 means for constructing a second encrypted string from the queue identifier and information specific to the customer terminal received as part of a subsequent request for service and the secret phrase; and

means for validating the subsequent request for service by comparing the first encrypted string with the second encrypted string.

25 Preferably, the system further includes means for checking at the service host that the request for service has been forwarded from the queue server. Preferably, the means for checking that the request for service has been forwarded from the queue server includes:

means for encrypting information specific to the request for service, information specific to the queue server and a secret phrase to form an encrypted queue server string;

30 means for sending the encrypted queue server string with the request for service to the service host; and

35 means for validating the request for service by comparing the encrypted queue server string with a service host encrypted string constructed at the service host using the secret phrase and information specific to the request for service and information specific to the queue server.

Preferably, the system includes means for automatically deleting the queue identifier and any encrypted string from the customer terminal following the provision of service by the service host.

Preferably, the system comprises means for automatically refusing requests for service after a predetermined number of requests for service have been forwarded to the service host.

Preferably, the system comprises a plurality of queue servers, and means for allocating requests for service from customer terminals to a particular queue server in accordance with a predetermined metric.

According to a third aspect of the present invention, there is provided a computer program product having computer executable code stored thereon for causing a computer to perform the method of the first aspect of the present invention.

Brief description of the Drawings

Examples of the present invention will now be described in detail with reference to accompanying drawings, in which:

Figure 1 is a graphical illustration of the behaviour of a typical server;

Figure 2 is a schematic representation of a typical e-commerce system in accordance with the prior art;

Figure 3 is a schematic representation of an e-commerce system in accordance with the present invention;

Figure 4 illustrates the relationship between queued customers, queue servers and a queue management system in an example of the present invention;

Figure 5 illustrates the stages of customer flow in the queuing system in accordance with the invention; and

Figure 6 illustrates an example of a screen-grab of a Queue page presented to the user's browser in accordance with the present invention.

Detailed description

Figure 2 illustrates a typical non-queued e-commerce system. The customer has a customer terminal 20. The customer terminal is shown three times in order to represent the three phases the user goes through during a purchase. The purchase depicted in Figure 2 consists of the following steps, indicated by arrows (a) - (o).

a) The user makes an HTTP (HyperText Transfer Protocol) request of a web server 21.

- b) The web server 21 responds with an HTTP response containing an HTML (Hyper Text Markup Language) page. The page is displayed on the user's screen.
- c) The user clicks on links to select other pages, or products to buy,
5 resulting in further requests.
- d) This results in further responses from the web server 21 containing different pages. If a user clicks on a link that indicates a product has been chosen, that information is stored on the server (not shown).
- e) Eventually, the user must click on a link that causes the server to
10 respond...
- f) ...with a page containing a "Buy" or "Check Out" button.
- g) When the user clicks on this button, the user's browser sends a request to a secure web server 22, usually protected by the HTTPs (secure) protocol. HTTPs is used to protect sensitive information such as credit card details. HTTPs has a load
15 penalty attached to both client and server machines due to the encryption, so HTTPs transactions are represented in the diagram with bold lines.
- h) The secure web server 22 responds with a page containing blank form elements for the user to enter his payment details, such as a credit card number.
- i) Once the user has filled in the form, and presses the "Submit" button,
20 the information is sent back to the secure web site. The pending order may at this point be saved to a database (not shown).
- j) The secure web site forwards the information to a payment gateway 23, again over HTTPs, and waits for a response.
- k) The payment gateway extracts the credit card number and amount and
25 other details entered, and sends them to a credit card network 24 for processing (this may or may not be over HTTPs).
- l) The credit card network responds with a decline, or an authorization for the amount requested. We shall assume the latter takes place.
- m) The payment gateway writes the order to a database 25.
- 30 n) The payment gateway sends the result to waiting logic 26 on the secure web server. The secure web server may write this to a database.
- o) The logic 26 on the secure web server sends an HTTPs response to the user indicating the order is successful. This is also usually followed up with an email (sent separately).
- 35 This is just one example of an e-commerce transaction. There are many variations. For example, steps (i) and (j) may bypass the Secure Web Site, so results

from the payment form are submitted directly to the payment gateway.

Steps (n) and (o) may bypass the Secure Web Site, so results from the payment gateway are submitted directly to the end user (this is rare, and requires further communication between the payment gateway and other systems to establish
5 that an order has taken place).

After the order has been placed (o), further interaction may occur between the various server systems, including the transfer of database information

The Web Site, Secure Web Site and Payment Gateway may all reside on separate machines. Or they may reside on the same machine. Alternatively, one may
10 reside on one physical machine, and the other two may reside on a different machine. Frequently, the Web Site and Secure Web Site reside on the same machine, but listen on different TCP/IP ports. They may still be considered as logically separate systems in most cases.

The type of system described in Figure 2 has no mechanism for dealing with
15 demand exceeding the system capacity in a satisfactory way. A system in accordance with the present invention incorporates an additional server 30 between a customer terminal and the secure Web server. The additional server 30 is a queue server which acts as another logical system in between the Web Site and the Secure Web Site. An example system in accordance with the present invention is shown in Figure 3.

Figure 3 is the same as Figure 2 except that step (g) of Figure 2 is replaced by
20 queueing steps (I)-(IV). The user terminal 20 is shown four times in order to represent the four stages of the transaction. The system includes a web server 21, secure web server 22, payment gateway 23 and credit card network 24, as described with reference to Figure 2. The new queueing steps are as follows:

25 I) When the user clicks on the "Buy" or "Check Out" button, instead of making a request of the Secure Web Site as previous, the user's browser makes a request of logic 31 on a Queue Server.

II) If the logic 31 determines that the user should be queued, the server
30 responds with a Queue page, containing a request number. The request number is stored on the user's machine.

III) The Queue page contains code that is executed on the customer machine, which determines when the user's request is likely to be due for service. This page also contains code that automatically refreshes the page with another HTTP request (III) at this time, or every few minutes, whichever is the sooner.

35 IV) If the logic 31 on the Queue Server determines that the user is at the front of the queue, the user's refresh request is forwarded to the Secure Web Site over

HTTPs, in a manner identical to (g) in the traditional setup. Otherwise, the server responds with a further queue page (back to step II).

As illustrated in Figures 2 and 3, the queue server of the present invention integrates into a known e-commerce system at a point immediately after a user has chosen the product, and clicks on the 'buy' button. A Target Page, e.g a credit card details page, is presented by a Target Service Host (in the Figures 2 and 3, this is the payment gateway and secure web server 22). In Figure 2, the user is taken directly to the Target Page; in Figure 3 however, the user is directed to a queue server 30, which in turn is connected to the credit card details page. In a preferred embodiment, the queue server is implemented using Java on a Tomcat J2EE Servlet Server. Queue servers have the following properties:

1. For incoming request rates less than the optimal rate of the payment gateway, the user is forwarded directly to the credit card details page.
2. For incoming request rates greater than the optimal rate of the payment gateway, the user is held in an automatically managed queue.
3. Users leave the queue on a first-come, first-served basis, at the optimal rate for the payment gateway (which may be some pages beyond the Target Page).
4. While they are part of the queue, users are informed of how long their wait is likely to be. When it is their turn, they are forwarded to the Target Page automatically.
5. Users may leave the queue at any time by closing the browser window (and optionally disconnecting from the internet). Their place in the queue is saved for them, and if they have missed their turn when they log back in, they are taken straight to the Target Page automatically. This is of particular benefit to users who connect to the internet through a modem, and may be subject to per-minute internet use charges.
6. Multiple queues may be employed simultaneously for different products, or to provide different access routes to a single product.

The system of the present invention as described, presents a number of benefits to both the providers of the e-commerce system (and the product being sold), and the customers using the system. Firstly, the risk of catastrophic failure is eliminated. Secondly, the owners of the e-commerce system experience a near-constant stream of customers at the optimal rate for them, which they can specify and control during the sale in real time. As they have regained control of how fast orders are received they can spread the sale over several days, substantially saving on hardware and staffing costs. Thirdly, the users of the e-commerce system perceive that they are part of a fair queue, and will be served on a first-come, first-served basis.

Even users who arrive too late to buy limited-quantity products (such as tickets) will perceive that they have been treated promptly and fairly by the system. This greatly increases user satisfaction. The users also have an indication of when they are likely to be served. This means that rather than continuously having to repeat their actions
5 in order to obtain the product, they can do something else with their time instead, and check back when it's their turn, resulting in much higher user satisfaction.

With a system in accordance with the present invention, users no longer repeat their actions, and as a result overall load on the system is greatly reduced, preventing back-propagation of catastrophic failure to the product web site (step 1), and further
10 reducing hardware costs.

The users experience a server that is responsive to them, increasing their satisfaction.

The queue management solution can also be used to prevent single individuals from buying large quantities of the product. This is particularly useful for vendors of
15 event tickets, who wish to reduce the number of sales made to ticket touts.

The queuing system of the present invention uses a numbered request system to handle new and returning users as follows:

1) On entering the queue, each new user is assigned a Request Number, which is stored on their computer as a persistent Cookie. The Request Number can also
20 be stored in an email, provided the customer enters their email address prior to joining the queue, or in the user's internet browser's Bookmarks or Favourites list, or simply encoded elsewhere in the web page response.

2) The system also keeps track of the number currently being served (the Service Number), which in a preferred implementation is automatically incremented at
25 a constant rate (the Increment Rate). As discussed below, this rate can be changed while the server is running to allow fine control of user pass through rates.

3) The response from the queuing system contains the user's request number, the service number, and the increment rate. The HTML response contains JavaScript code that runs on the user's computer to transform this information into a visual
30 display, which may contain one or more of the following elements.

- The user's request number,
- The service number,
- The pass through rate,
- The number of people ahead of the user in the queue (given by the request
35 number minus the service number),
- The estimated wait time (given by the number of people ahead in the queue

divided by the pass through rate). This may be displayed in a count-down format,
-Advertising (in standard internet Banner format) or other content from third
party sources.

Figure 6 illustrates an example of a "Queue page" visual display presented to
the user's browser and showing both the user's request number and the service
number.

In addition, the code running on the user's computer automatically refreshes
(reloads) the page every five to ten minutes, or when the wait time reaches zero,
whichever is the sooner.

When the user makes a further request to the queue server, either by a manual
reload or through the automatic refresh mechanism, the previously assigned request
number is automatically sent to the server as a cookie along with the request.

If cookies are not being used, the request number can also be sent as part of
a link in an email, in the user's Bookmarks (i.e. the user's browser's favourites list), or
on a web page (as a Query String, e.g. ?id=5323 within a URL, or HTTP GET
parameter), or as a stored Form element inside a web page or HTML email (which may
be sent back as an HTTP GET or POST parameter). Further alternative storage
mechanisms include storage of the request number as a variable within JavaScript, or
some other element of the HTML page, and storage in a dedicated database.

The queue server compares this number to the current service number. If the
request number is less than or equal to the service number, the user is passed through
to the Target Page. This allows users to return after their number has been 'called',
yet still complete the transaction (assuming there are quantities of product left to buy).

If, however, the request number is greater than the service number at the time
the request is received, the user is shown a similar page containing the new service
number and estimated wait time, and/or number of people ahead in the queue, which
may vary if the pass-through rate has changed. No new request number is issued in
this case.

Each queue consists conceptually of two components as follows:

1. A single persistent Assigner component, responsible for keeping track
of the last request number issued, the automatic increment rate, and the current
service number.

2. Multiple, ephemeral, identical Request components, one for each
request from a user, responsible for getting new request numbers from the shared
Assigner component if necessary, for checking incoming request numbers from
returning users against the current service number, and for returning the appropriate

response to the user.

This system has a number of technical advantages, which combine to allow it to cope with very large queues. The system is fully automatic, in that no data need be entered manually by the users in order to participate in the queue. No per-user data is stored on the server (individual request numbers are only stored on the client machines), negating the need for per-user database writes, and further increasing availability. No sensitive data is transferred between the user and the queue system, negating the need to use HTTPs as the transport protocol for the system, further increasing availability. Although unnecessary, it should be noted that the queue system of the invention can be further protected by HTTPs where this is required. The work performed by the server is minimized, in that presentation logic is performed by the client machine, allowing for even greater response rates.

Typically, a user accessing the queue system will already have chosen a product, but will not yet have entered in any personally identifiable information.

When the user leaves the queue and is sent to the Target Page, the Target Service Host will typically need information to identify these choices in order to correctly process the user's order.

The invention addresses this need by automatically converting any HTTP GET (query string) or POST parameters sent to the server when the user is initially forwarded to the queue into cookies, and storing them on the user's machine. These parameters can also be stored inside a web page, or email as before. When the user is eligible for service, these parameters are then forwarded on to the Target Service Host. The designers of the protected systems (the Target Service Host systems) are free to choose exactly what information is stored to meet their requirements.

In one example, the choice information itself may be stored on the user's computer by the Queue Server. Alternatively, the choice information may be stored by the Target Service Host, and a customer identifier may be stored by the Queue Server on the user's computer so the Target Service Host may later recover it.

This latter approach is commonly used when session technology is used to store user data, however Target Service Hosts relying solely on the storage of a session identifier must be careful to ensure that the session does not expire before users leave the queue.

While the solution discussed above sends all queued users to a single Target Page, the system is also capable of sending queued users to a range of Target Pages, possibly on different Target Service Hosts. This may be useful in a number of scenarios.

In one scenario, a single queue may be used for multiple product types. A typical instance being the sale of a range of different ticket types (for instance UK-resident and international tickets). While it is certainly possible to encapsulate this choice using the stored GET or POST parameters of the previous section, it is also possible to implement this by sending the users to different target pages.

Another situation where the provision of a plurality of Target Pages is useful is if an event organiser wishes to sell tickets through multiple web-site outlets. Users may feed into a single queue from each of these web sites, and then be forwarded back to the corresponding web site when their turn comes up. This allows the event organizer to ensure first-come, first-served fairness across multiple points of sale, allows simultaneous protection of all the vendor sites, and allows precise control of exactly when tickets become available on each of the sites.

The system in accordance with the invention can also be configured to use a plurality of Target Pages to load balance users leaving the queue automatically across several different servers, if desired, using a variety of load-balancing algorithms.

It is desirable that the system be responsive to administrative requests to manage the queue (for instance, to lower the Increment Rate if the Target Service Host exceeds its optimal rate), and also that the queue be recovered in the event of a server restart due to unforeseen problems. This is achieved as follows.

All the individual Request Numbers are stored on the client machines, so server restart does not disrupt an individual's position in the queue. The current Request Numbers and Service Numbers are written to persistent storage on a regular basis (for example, once every ten seconds), to allow automatic recovery of the queue in the event of a server restart or database malfunction.

The current Request Numbers, Service Numbers and other queue parameters are also checked against input data on the persistent storage (e.g. the internal file system and/or an external database) once every ten seconds, allowing control of these values while the server is running.

The input data (and persistent storage) is administered remotely through a separate, dedicated internet connection, ensuring that the server can be controlled even under very high loads.

There are a number of Increment Rate strategies that are appropriate in different situations. The automatic increment of the Service Number (which dictates which queued customers are eligible for service) has already been mentioned. In the simplest implementation of the inventive system, this Increment Rate is a constant (expressed in users per minute), which can be reset on the fly by system

administrators.

Constant Increment Rate is a good strategy when it is anticipated that customers will attempt to use the service very soon after their Request Number has been 'called'. However this may not always be the case; for instance in the case of a queue that is going all night, people may avoid purchasing in the small hours, even if their turn comes up. This may result in a glut of customers in the early morning, seeing as customers usually pass through the system any time after their Request Number has been called.

A more sophisticated strategy is to monitor the number of actual people passing through the queuing system in a particular minute, and reduce (or increase) the increment rate if this is more than (or less than) the optimal rate of the Target Service Host. In extreme cases, this may lead to the Service Number actually decreasing (as opposed to increasing more slowly).

Another strategy is to have the Target Service Host control the Increment Rate directly through the use of messages sent from the target service host to the queue servers (according to some metric calculated by the Target Service Host). This is usually to be avoided due to the additional work performed by the Target Service Host, but can easily be implemented if required.

Depending on the size of the incoming spike, many queue servers may be required to deal with all the requests from people in the queue. A typical queue server in accordance with the invention operates at a rate of 30,000 per minute per server. As the server needs no connections to an external processing network or database in order to function, additional servers may be added to allow for arbitrarily large queueing demands.

In this case, each queue is replicated on each server. Each server maintains its own Assigner component, with its own current request number and service numbers. The servers share the same input data and database for control of queue parameters, so changes made to them affect all servers within ten seconds. The servers synchronize with each other once every ten seconds in order to provide consistent numbers to users.

Users can expect consistent results from each server, so the servers can be assigned at random to incoming requests, regardless of which servers a user may have interacted with before. This negates the need for 'sticky' load balancing between the servers, in which a particular user is always directed to the same server, further reducing hardware costs.

Figure 4 illustrates load balancing when multiple queue servers are used. Also

illustrated is a management system. A queue administrator terminal 40 is connected to a management system 41. The management system 41 is connected to each of the parallel queue servers 42. Queues are set up by the queue administrator. The queue administrator 40 interacts with the queue management system 41 at step (a). The queue administrator uploads template files and other data through HTTP requests and responses to and from the management system 41. The management system 41 stores the templates and data in its own persistent storage (consisting of a database and/or file system as appropriate). The data held by the management system 41 is propagated to the memory and file systems of the queue servers 42 at step (b).

In one embodiment, each queue server 42 holds data relating to a number of different queues created in this manner. Once the queue servers 42 have been suitably configured by the management system, customer requests can be queued on the queue servers via a load balancer 43. Communication between the load balancer 43 and the queued customers is indicated as step (c). The load balancer 43 is responsible for distributing requests from the users evenly across the configured servers, in order to make best use of the available bandwidth. No per-user data is stored on the system and so any server may serve the queued user and sticky or session based load balancing is not required. Once the request is received at a queue server it is assigned to the appropriate queue and given a Request Number. The queuing process then proceeds as described with reference to Figure 3.

In cases where more than one server is used, the servers can also be configured to increment the issued Request Numbers in intervals other than one. For instance, if two servers are used, one may issue Request Numbers in the sequence (1, 3, 5, 9...), and the other in the sequence (2,4,6,8...). This ensures that no two customers share the same Request Number.

Optionally, the Assigner components may be configured to obtain new Request Numbers from a single source to ensure first-come, first-served fairness, though this may impose limits on the number of new people queued per minute.

Separate queues (domains) may be established, based on customer or product identifying data. Where appropriate, the domains themselves may be used to balance load between servers. This separation of queuing is referred to as Domain Queuing, and is best explained with an example.

Consider an event that wishes to sell tickets across the nation. The event is likely to be heavily oversubscribed, so the queue system is deployed. Due to the high interest, the queue becomes full (that is, has as many queued customers as tickets available) in a very short time (say 5 minutes).

If only one queue is used, people with better internet access (e.g. lower latency, higher effective bandwidth, better contention ratio) are more likely to gain access to the queue servers during this time due to network saturation. This means a disproportionate amount of tickets will go to people living in major cities, which are
5 closer to the internet backbone for the country.

One way of solving this is to employ multiple queues based on geographic customer data. For instance, if the customers enter the first letters of their postcode (e.g. 'EX' for Exeter) before entering the queue, a separate queue can be deployed for each of these geographic domains.

10 This ensures that each region is queued separately. Though it may take customers from remote regions longer to reach the queue server, they are still queued in order and with equal priority to those from major cities.

Other 'domains' may be used. For instance, customers may be allocated to separate queues based on demographic (i.e. age, gender etc.) information, or on
15 product information such as price point.

Where possible, domain information should be checked on purchase of product to prevent fraud (for instance, the supplied postcode should be checked against that entered in the Credit Card Form when tickets are purchased in the above example).

Each of these queues is handled by a separate Assigner object within the
20 framework of the queue system. It is also possible to allocate a certain amount of product to each of these queues, and to shut down any particular domain queue automatically when this amount has been sold. This gives product providers fine control over the distribution of product even in very high load situations.

Furthermore, it is possible to control all these queues using a single Service
25 Number for simplicity.

As a specific example, consider a situation where an event wishes to sell 100,000 tickets. 90,000 of these are to be distributed to UK residents (domain 'UK'), with the rest going to international customers (domain 'Int'). In this case, the queue servers use floating point arithmetic to allocate request numbers in the sequence 10,
30 20, 30, 40,... to the 'Int' queue, and in the sequence 1,2,3...9,11... to the 'UK' queue. This means that when the (single) Service Number reaches 100,000, there are 90,000 UK residents eligible for tickets, and 10,000 international customers eligible, as required.

The queue system is conveniently integrated with the sales system it protects.
35 This arrangement serves both to prevent fraud (which is discussed below), and to enable automatic management of queues based on sales events. A clear illustration

of the benefits of an integrated queue system can be seen where all tickets have been sold for a particular domain, people left in the queue should be told of this so that they may leave the queue.

5 Sales Tracking data can also be used to provide an indication to the Site Administrators of how the sale is progressing, and whether they need to transfer tickets from domain to domain in order to sell out. This can be done on the fly during the sale using the queuing system of the invention.

10 The inventive queuing system supports three ways of closing queues based on sales data. Firstly, it may simply use manual queue closure. In this method, site administrators for the Target Service Host use a web interface to indicate manually that all tickets have been sold. They may do this on a domain-by-domain basis, or may simply declare that all tickets have been sold for all domains. In either case, queued users are shown a different page indicating that the sale has finished on their next visit to the queue servers, and new users are not allocated new Request Numbers. 15 Depending upon the desired configuration, SMS or email messages may also be sent to users who have entered the required information.

Alternatively, queuing system may use integrated automatic queue closure. In this method, buying ticket(s) at the Target Service Host causes a message to be sent to the queue servers indicating that a certain amount of tickets have been sold. 20 The message may be sent directly from the Target Service Host, or may be sent from the user's browser through an embedded request to a file on the queue servers inside the 'Order Successful' web page that the user is shown after being allocated tickets.

In both of these cases, the message should include the number of tickets bought, any domain information required, and the purchaser's Request Number (to 25 prevent duplicate messages being processed as subsequent orders).

The queue servers must be preconfigured with a maximum number of tickets to sell for each domain. When the indicated number of tickets sold for a particular domain reaches this maximum, the queue for the domain is closed as above.

30 In a further alternative way of closing queues based on sales data, a simple automatic queue closure may be used. In this method, there are no messages sent from the Target Service Host. Instead, the queue server closes the queue for a particular domain after a prespecified number of people have passed through it.

There are a number of types of fraud that should be prevented in queue-managed e-commerce systems. These are:

35 1. Queue Jumping. This is where a user is already part of the queueing system, but attempts to reduce the value of his/her request number in order to be

served faster, possibly with the co-operation of another user ahead of him/her in the queue.

2. Queue Bypass. This is where a user attempts to reach the payment page without having been through the queue management system at all.

5 3. Multiple purchase. This is where a user, having been through the queue system, makes multiple purchases of the product. This may only be undesirable in some circumstances, such as a ticket sale as discussed above.

4. Domain Fraud. This is where a user attempts to buy product outside their stated domain.

10 While no computer system is ever completely free from the potential for misuse, the present invention may include steps to prevent these frauds.

The queue jumping problem is addressed as follows. Firstly, the use of cookies to store the request numbers represents a substantial barrier to the casual internet user, who would need considerable technical savvy to modify their contents. In order
15 to prevent fraud by people who have such savvy, an additional system is used, based on hashing functions.

A hashing function is a function that takes a lexical string as input, such as "Hello World!", and produces another lexical string as output, such as "AH74KFEu". Hashing functions used for security processes produce output that is sensitively
20 dependent on the input. For instance, the hash of the slightly different string "Hellq World!" should be something completely different, such as "Klo9SydP". Furthermore, there should be no way to recover the original string from its hashed counterpart.

The queue system of the present invention can use the industry-standard MD5 hashing algorithm to implement fraud protection as follows. When a user is initially
25 assigned a request number, for instance 1537, this is combined at the server with a sale-wide secret phrase to produce a lexical string, for instance "MyEventTickets1537". This string is prepended with an additional string that (ideally) uniquely identifies that user's computer. A preferred implementation uses the first two bytes of the user's IP address, to give a final string of, for example, "209.157MyEventTickets1537". This
30 string is then hashed by the server, and the hash is stored along with the original request number on the users's machine as a cookie. The secret phrase is never transmitted to the user's computer directly.

The stored hash string is now used to validate request numbers when the user returns to the queue server as follows. When a user returns to the queue server, the
35 original request number and the stored hash string are both returned to the server as part of the page request. The server then reconstructs the proper hash string for the

returned request number from the secret phrase and IP address, and compares it to the hash string returned by the users computer.

If the two strings are identical, the user's request is considered valid, and handled appropriately.

5 If the two strings differ in any way, the user's request is considered invalid. The user is issued a new, higher request number and hash string, and is in effect sent to the back of the queue as punishment for cheating.

10 It is important to try to uniquely identify the user's computer in the input to the hash function, as otherwise request-hash cookie pairs could be shared to allow many individuals to use the same request number, effectively allowing them to jump the queue. The implementation described above uses the first two bytes of the IP address because many people do not have static (constant) IP addresses, particularly if they log off the internet and come back later. However, people with dynamic (changing) IP addresses usually have the same first two bytes due to the technical implementation
15 adopted by their Internet Service Provider.

 Further uniquely identifying data can also be found in the browser identification string that is a part of every HTTP request.

 Additional fraud protection may be obtained by taking personal information, such as the last four numbers of the credit card or a postcode associated with the card,
20 that will be used to make the purchase before joining the queue, and adding this to the hashed string.

 Having identified a way of validating users in the queue, it is possible to also use this scheme to prevent users bypassing the queue as follows. The target page of the queue (usually the credit card details page) checks that incoming users have been
25 referred from the queuing page using the HTTP request referrer field. All other users are immediately sent to the queue for processing.

 The target page also performs the same validation check that the queuing system performs using hashes of the request number. Users failing the check are sent back to the queue for processing.

30 If desired, the target page may perform a validation check using a different secret phrase, to distinguish between users who have been through the queue, and users who may still be in the queue.

 In order to prevent individuals from making repeated purchase transactions, possibly with multiple credit cards, the payment gateway response page that indicates
35 a successful transaction to the user should delete the stored cookies representing the original request and security hash.

Any further attempt to buy the product would then automatically result in the user being sent to the back of the queue. In heavily oversubscribed sales, this will effectively prevent them from making multiple purchases.

Where Domain Queuing is used, a customer may attempt to purchase product
5 outside their stated domain. For instance, a person in London, on realizing that the London queue is full, may instead report that they are from some other domain. The queue system stores all domain choices (in this case the entered postcode), and forwards them to the Target Service Host. The Target Service Host should ensure that these match any subsequent order information supplied by the customer. In this
10 example, the supplied postcode should match the postcode entered in the credit card form, or delivery address.

Optionally, Domain data may also be included in the string used for the security hash, or stored in the queue database for later checking when other points of sale are used.

15 It has already been remarked that the storage of queue identifiers allows customers to turn off their computers while they are in the queue without losing their place. There is therefore no guarantee that a particular customer will be online when their turn comes round, particularly if the queue is likely to be in place for an extended period of time.

20 It is therefore desirable that customers can be notified when the Service Number exceeds their Request Number. To do this, the Queue servers can be configured to store customer communications data, such as an email address and/or phone number, and to send an email or SMS text message when appropriate. The Queue system does this by storing the email address and/or phone number in a
25 database, along with the request number, and periodically checking this database table against the current Service Number and issuing the required messages automatically.

While the system can be configured to obtain this information from all users before entering the queue, this is not recommended for performance reasons. Rather, it is better to first allocate each user their request number, and then provide them with
30 the (optional) means to register their email address or phone number for notification at their request.

In many cases, an initial sale of product will sell out, and then additional product may become available later. An example would be an airline flight. Currently, airlines sell more tickets than there are seats on any particular flight, as they expect a certain
35 number of cancellations, and must ensure each flight is filled to capacity. This leads to customers being 'bumped' to other flights when the expected number of

cancellations is not reached, leading to additional costs and customer dissatisfaction.

With the queueing system of the invention, the airline can suspend the queue when all the seats on the flight have been sold. After that, each new potential customer joins the queue in order, and supplies contact information for Customer
5 Notification. As cancellations are made, the cancelled tickets may be offered to the queued customers in order, by increasing the Service Number at a slow, constant rate, until the cancelled tickets are sold.

In the light of the foregoing, it should be clear that present invention is applicable not only to e-commerce but to any scenario in which demand over a
10 communications network might exceed capacity. It should also be made clear that when using the Internet, the requests for service in the method and system of the present invention can use either HTTP GET parameters or HTTP POST parameters, whichever are the most appropriate in the circumstances.

The method of the present invention may be encoded as software which runs
15 on existing hardware devices. Alternatively, specific hardware or firmware may be built to implement the invention. The present invention may be provided as use of a queue server remotely over a communications network as a service to service providers. Alternatively, a suitably programmed queue server may be integrated into a service provider's existing architecture.

20 As will be readily apparent, most ticket sales do not take place solely online. Usually, customers are also able to call a phone line. Sometimes other points of sale, such as physical box offices or SMS text devices are also used. The invention facilitates the integration of telephone, internet, and other queue systems, thereby allowing customers who join the queue at one point of sale to purchase product at
25 another.

In an equally important embodiment, the present invention may also be applied to telephone product sales systems. The telephone embodiment has similarities to the internet (e-commerce) embodiment described above. There is, however, one major difference: telephones are not, in general, able to store data, so Request Numbers
30 must be stored at the server.

Situations providing sales or other services provided by phone are frequently oversubscribed. This may result in a queue of customers 'on hold' due to a shortage of operators, or new customers receiving the engaged tone if the service provider has run out of phone lines to connect them. Both cases result in customer dissatisfaction,
35 and may result in additional telecommunications expenses to both the service provider and the customer.

Rather than presenting a HTML front end, a Telephone queue system in accordance with the invention is provided with an automatically generated audio message. The audio message(at least) indicates that the customer is in a queue, and updates their expected wait time on a regular basis.

5 There are two slightly different implementations of the telephone queue system embodiment, depending on whether multiple places in the queue are to be supported for a single telephone number.

 In the simplest telephone system, each telephone number is associated with just one Request Number (one place in the queue). Customers calling the telephone
10 line are routed through a Telephone queue server (using, for example, the Asterisk open source PBX). The queue server first checks the incoming telephone number using a Caller ID function. If Caller ID information is unavailable, the queue server can be configured to either ask the caller to hang up, enable Caller ID, and try again, or prompt the user to enter their telephone number using the numeric keypad.

15 Optionally, the Telephone queue server may ask the customer to provide the last four numbers of their credit card, or other personal information, to provide additional fraud protection.

 The queue server then checks the telephone number against a database of telephone numbers that are already in the queue.

20 If the number is not in the database, the caller is 'new', and the Telephone queue server issues a new Request Number, and stores it along with the telephone number in the database. Optionally, Domain Queuing may also be used, using voice recognition or data from the numeric keypad to specify domain information before the Request Number is issued.

25 If, on the other hand, the telephone number is already in the database, the caller is 'recognized', and the corresponding Request Number is retrieved from the database.

 If the Request Number is greater than the current Service Number and if the customer is new, the new Request Number is read to the caller, and the system may
30 also send it to the customer using an SMS text. The caller is then told the current Service Number, and given an estimate of how long the wait will be to become eligible for service.

 If, alternatively, the Request Number is greater than the current Service Number but the customer is recognized, the caller is told that their phone number has
35 been recognized, and may be reread their Request Number. The customer is then told the current Service Number, and given an estimate of how long the wait will be to

become eligible for service.

Callers are given the option to hang up, and call back after a personalized, pre-arranged time, when they can expect to be connected immediately. An example message is: "Hello. Due to high demand, you have been placed in a queue. Your wait
5 time is 10 minutes, 24 seconds. Your phone number has been stored by our system and will be recognized, so please feel free to hang up and call back then, when you will be served immediately." This saves the callers the expense of being on hold.

Optionally, the system may hang up at this point. If the service is running out of phone lines to handle new customers, the customer is given the following message:
10 "Hello. Due to high demand, you have been placed in a queue. Your wait time is 10 minutes, 24 seconds. Your phone number has been stored and will be recognized, so please call back then and you will be served immediately. This call will now disconnect." This saves the callers the expense of being on hold, and saves the provider the expense of running multiple phone lines to keep them on hold.

15 The system may also automatically call back customers on the stored telephone numbers when it is their turn. Instead of forcing users to call the service back, the service automatically calls them back when an operator becomes available. This saves the caller the expense of a second phone call (which instead is borne by the provider).

20 If, on the other hand, the Request Number is less than the current Service Number, the customer is eligible for service. The customer's call is either serviced by the Telephone queue server, or is forwarded to a Target Service Host (another PBX in this telephony system), or to an operator for service.

Rather than tell the customer their Request Number or the current Service
25 Number, the system may instead simply give an estimated wait time.

Optionally, the Target Service Host (the PBX through which the product/service is available) may cause the Telephone queue system to bar telephone numbers, or remove them from the database, after product has been sold to prevent fraud (e.g. by ticket touts).

30 In a telephone queue system, multiple places in the queue may be allocated to a single telephone number. For instance, one might have a case where an entire office of people uses a single phone number to make outgoing calls. The queue system first obtains a telephone number as above, and compares it to the list of telephone numbers stored in the database.

35 If the telephone number is not in the database, the telephone is considered 'new', otherwise the telephone is considered 'recognized'.

If the telephone number is 'new', a new Request Number is issued as in the one to one mapping case above, and a numeric Confirmation Code is also generated using a hashing function of the telephone number, the Request Number, and a secret phrase as before in the Internet solution. The Confirmation Code is stored along with the Request Number and telephone number in the database. The Confirmation Code is read to the caller, and may also be sent in an SMS text as before. The caller is considered 'new'.

If the telephone number is 'recognized', the caller is asked whether they are already in the queue, and answers by pressing the numeric keypad on their phone or through voice recognition technology.

If the caller indicates that they are already in the queue, the caller is asked to enter a Confirmation Code. If this matches a stored entry in the database with a matching telephone number, the caller is treated as 'recognized'. If not, the caller is prompted to re-enter their Confirmation Code, or (optionally) assigned a new Request Number and Confirmation Code as follows.

If the telephone number is 'recognized', but the caller indicates that they are not already in the queue, a new Request Number and Confirmation Code are generated, sent and stored as above. The caller is considered 'new'.

New and recognized callers are then treated as described in the one to one mapping case above.

After product has been sold to the caller, the caller may be barred or deleted from the database by removing or altering the record corresponding to their telephone number and Confirmation Code.

In this way, multiple customers may use the same telephone number, but be treated as individuals by the system.

In a variation on the Telephone queue system embodiments, the Telephone queue system can be easily modified to provide a queuing service over SMS (so-called 'text messaging').

As in the Telephone queue system, where there is a one to one mapping embodiment, the customer enters the queue for the first time by sending a text message to a SMS queue server. The text message may also contain Domain information used to route the customer to the appropriate queue. The server responds with a text message back to the customer's telephone (optionally) indicating the customer's Request Number, (optionally) the current Service Number, and an estimate of how long they will be waiting. The Request Number and (mobile) telephone number are stored in a database as in the Telephone queue service embodiments.

Subsequent text messages sent to the SMS queue server result in a response indicating an updated wait time.

When the Service Number exceeds the issued Request Number, the server may respond by purchasing product automatically (and billing the user through the user's telephone service provider), or by notifying the customer to take some other action, or by allowing the customer to take some other action.

An SMS queue service embodiment of the invention with many to one mapping is also provided. This embodiment is identical to the one to one mapping case above, except that the message returned on entering the queue contains a Confirmation Code. This code must be returned in subsequent messages by the customer in order for the system to distinguish between different customers using the same phone.

Another embodiment integrates the Internet and SMS systems described above, using email facilities. Most modern email clients are capable of displaying HTML email pages, though some are not. Some email clients are also capable of storing cookies.

Customers enter the queue by entering their email address on a web page, or by sending an email to a particular email address. The email or web page may gather other data for Domain queuing. An (email) queue server responds by generating a Request Number, and sending an email back to the received address.

The response email may contain a link containing the user's request number and a security hash (though there may be no way to ensure that this security hash is constructed from terminal-specific data, the customer's email address may be used for this purpose).

This link can then be clicked on, or pasted into a web browser to subsequently join an internet queue. The server responds to the request encapsulated by the link by storing the appropriate queue information on the customer's terminal, or at the queue server as appropriate. The customer is then part of an Internet queue.

The email may also contain an element (such as a graphic) downloaded from the queue server; this download may also cause the user's browser to store necessary information to participate in the internet queue, without the customer having to paste or click on the link through the use of persistent cookies or the user's bookmarks.

The use of email to send credit card data is discouraged, as email is not a secure method of information transfer, so users entering the queue by email are likely to complete their purchase using another point of sale.

Optionally, the email may also contain a Confirmation Code that the customer may enter into the telephone based system.

For convenience, embodiments of the system can be used in physical retailers (box offices, stores etc.). In this system it is envisaged that the retail staff have access to the internet or telephone, and enter customer's details into the appropriate version of the inventive queuing system on their behalf.

5 In addition to the Request Number, a Confirmation Code may be issued by the web server, and given to the customer, who must use the code in addition to the issued Request Number to access their place in the queue.

10 If possible, the last four digits of the customer's credit card, or other personal information should be used to generate the Confirmation Code, so that it can be checked later.

The Request Number and Confirmation Code may be used to access their place in the queue through other points of sale if the customer chooses not to return to the retailer in order to buy product (this scenario is discussed further below).

15 On the other hand, if customers do wish to complete the purchase through the retailer, they would return with the Confirmation Code and Request Number, and any personal information supplied. The retailer can then complete the sale for them through the queuing system, or may complete the sale using existing point of sale equipment. In the latter case, the queuing system should be notified of the sale in order to facilitate automatic queue closure (thus, for example, all Request Numbers, Telephone Numbers and Confirmation Codes will be sent to an Internet Queue Server).

20 In sales where the rate of sale becomes less than the Increment Rate for a certain period of time, such that new customers should pass through the queue immediately and transparently, the queuing system can notify the retailers of this, so that instead of helping new customers enter the queue, they simply sell the product.

25 Note that if cash is used to buy tickets, then other personal information is needed for fraud prevention, if desired.

30 The scenarios described above are not the only instances where a customer may enter the queue using one point of sale, and purchase tickets, services etc. using another point of sale. Specific examples include queuing on internet, and buying on telephone; queuing on telephone, buying on internet; and queuing on email, buying through other means.

Note that if an integrated solution is adopted, the use of Confirmation Codes is encouraged to prevent people from fraudulently entering someone else's telephone number into the system.

35 In all integrated cases, once the customer has bought product through one point of sale, the servers handling the queue for other points of sale may be notified to

prevent further purchasing (for instance to prevent multiple sale by ticket touts).

In the scenario where customers queue on internet, and buy over the telephone, a customer enters the queue online using the system already discussed. On the queue page, the customer is given the option to buy tickets by telephone. The
5 customer enters their telephone number, and this is stored along with their request number, and transmitted to the telephone queue server.

If many to one mapping is used on the telephone queue server, then a Confirmation Code must also be generated and given to the customer, and sent to the telephone queue server.

10 When the customer then calls the telephone hotline, the necessary data is already present in the telephone queue server's database, and the customer is treated as 'recognized'.

If an SMS point of sale is also used, the internet queue server need only send the information to the SMS queue server for the customer to be similarly handled.

15 In the scenario where the customer queues on the telephone and buys on the internet, he enters the queue using the telephone queue system as discussed, and is given a Request Number and (optionally) a Confirmation Code. The telephone number, Request Number and (where present) Confirmation Code are sent to the Internet Queue Server, where they are stored in a database. The customer then logs
20 onto the internet and reaches the Internet Queue Server.

As the Internet Queue Server has no means of recognising this customer automatically, the customer is issued a second Request Number, and is shown the Queue Page. Because they entered the telephone queue first, this Request Number will usually be larger than the original Request Number obtained by phone.

25 On the Queue Page is an option for people who have already joined the queue by phone to use their telephone-issued Request Number. Customers selecting this enter the Request Number they received by phone, their telephone number, and (where present) their Confirmation Code. If a matching record is found, any Request Number or other queue information stored on the customer's terminal is deleted, and
30 overwritten with the information found from the database. The customer can then proceed through an internet queue as before.

Alternatively, customers reaching the internet after calling the telephone system may be given the option to enter their telephone number before being issued an internet request number.

35 As the SMS queue server also provides the customer with identical information to the Telephone queue server, the same mechanism can also allow customers who

have joined the queue by Telephone to buy their tickets by SMS (and vice versa), or for customers who have joined the queue by SMS to purchase product online (and vice versa).

Another example where the queue is entered in one medium and the point of sale is another is where queuing is on email and buying is through other means. Email is, as stated above, not a secure means of transmitting credit card information. Therefore, some other means must be used. Customers entering the queue by email can be directed to arrive at the internet queue; from there they may proceed to any of the other points of sale.

Figure 5 shows one possible configuration of customer flow in the queuing system of the invention. The Figure conforms to the well-known UML (Unified Modelling Language) Activity Diagram specification. In this representation technique, Initial States (New Customer, Returning Customer) are marked as solid black dots. Final States are marked as solid black dots with a circle round them (Queued Customer, Order Complete). Actions and States are represented as rounded rectangles. Decisions are represented as diamonds, and flows (on completion of Actions or results of Decisions) are represented by arrows. Finally, dashed arrows represent optional user actions.

The system is divided into three "Swim Lines": Customer; Queue Server; and Payment System, which represent different actors (i.e. components) of the overall system. The Figure applies equally to both the Internet and Telephone embodiments of the queuing system described above, as will be appreciated from the following description.

For the Internet system, Customer represents: the Customer; the Customer's browser; and the web site through which the Customer makes product choices. Payment System represents a Payment Gateway responsible for taking and processing credit card information. Queue Server represents the Internet Queue Server.

For the Telephone system, Customer represents the Customer and the Customer's telephone. Payment System corresponds to a call centre, its operators and associated software. Queue Server represents the Telephone Queue Server.

New Customers 502 enter the system and choose a product 510, either through a web site, or by calling a phoneline, or in some cases by sending a text or email. Optionally, they may be required to enter additional personal information 512, which may be used as later security information (for instance, the last four numbers of the credit card they will use to complete the purchase), or for some other purpose.

Once this information is gathered, Customers are forwarded to the Queue Server 520. As the Customer is new, the product choice and any other information gathered is stored 524 (in a database in the Queue Server in the case of the Telephone system, or on the user's machine as a cookie in the case of the Internet system). A new Request Number is generated 526, and stored, and is compared to the current Service Number 528.

If the Request Number is less than the Service Number, the Customer is forwarded to the Payment System for processing 546.

If the Request Number is greater than the Service Number, then the Customer is given queue information 534. In the Internet case, this is typically a webpage showing the Customer's Request Number, the current Service Number and an estimated wait time. In the Telephone case, this information is read to the Customer. The Customer is now queued.

This response may also include information allowing the Customer to return to the queuing system from another point of sale, and be recognized 536. This varies by system:

- for the Telephone system, Recognition Information is usually the Customer's telephone number, and may also include the Customer's Request Number, and optionally an automatically generated Confirmation Code.
- for the Internet system, usually no Recognition Information is sent, but the Customer may be asked to enter their telephone number if they wish to buy product by phone later. This Recognition Information is then stored.

The Customer may optionally request to be notified when they reach the front of the queue (specifically, when their Request Number becomes less than the Service Number). In this case, the Customer enters their telephone number and email address (internet system only; on the phone system the telephone number is already known) 538. The queue system waits for the Customer's Request Number to become eligible for service 540, and sends an SMS text message and/or email to the Customer 542.

Notification messages may also be sent if additional product becomes available at a later date (not shown).

Customers returning to the system 504 as a result of receiving one of these messages, or a subsequent phone call, or an automatic page refresh, are represented by the Initial State marked as Returning Customer.

Returning Customers may arrive through the same point of sale (i.e. they have already joined the queue by Internet, and are returning over Internet), or may arrive through a different point of sale (i.e. joined the queue by Internet, and are returning by

calling the helpline).

Depending on the specific points of sale involved 514, the Customer may need to enter Recognition Information in order for the system to retrieve the Request Number the Customer received earlier 516. In many cases this Recognition
5 Information, and the associated Request Number, may be retrieved automatically.

Once the Request Number has been retrieved 530, it is checked for validity as a fraud prevention measure 532. If the Customer's information fails the validity check, the Customer is issued a new Request Number 526, and is effectively sent to the back of the queue.

10 If the Request Number is retrieved and successfully validated, the Request Number is compared to the current Service Number 528. If the Request Number is still greater than the Service Number, the Customer is sent updated queue information 534, including an updated wait time.

If the Request Number is now less than the Service Number, the Customer is
15 sent to the Payment System for processing 546. All gathered information is also sent to the Payment System, where it may be stored or otherwise processed.

The Payment System may use this information to check that the arriving Customer has passed through the queue system 548. This is particularly important on the Internet, where care must be taken to prevent users from arriving at the Payment
20 System directly. The Payment System may perform this check itself, or may send a message to the queue server requesting validation (not shown).

Once the Customer is successfully validated, the Payment System asks the Customer for credit card details 550. These may be checked against Personal Information gathered by the system before entering the queue to provide further fraud
25 protection (not shown).

Assuming the credit card information is successfully processed, the Queue Server may be notified that the purchase has taken place 556, in order to prevent the same Customer purchasing product multiple times (particularly important in ticket tout prevention). The order is then complete 558.

30 As will be apparent to the reader, various stages of the customer flow described above may be omitted, replaced or augmented. Furthermore, there is flexibility in the "actor" at which many of the stages of the flowchart are carried out, and the order in which the stages are executed.

In the models so far described, it has assumed that each point of sale, such as
35 telephone or internet, represents a separate domain; that is there is a separate allocation of tickets to each point of sale, and that Request Numbers and Service

Numbers increment for each separately.

These assumptions are reasonable, and may indeed be necessary in the case of a heavily oversubscribed sale, in which the Internet Queue is likely to fill very much more quickly than the Telephone Queue.

5 Tighter integration may be achieved in a number of areas, in particular: Service Numbers; Request Numbers; and Domain Switching.

 The Service Number indicates which customers are eligible to buy tickets. Typically, the Service Number is incremented at the maximum rate that the Target Service Host can sell tickets. This is likely to be faster for the Internet than the
10 Telephone, particularly when human phone operators are taking the payment details.

 In cases where a single Service Number for all points of sale is required, this should be chosen to increment at the rate of the slowest Target Service Host (usually the telephone).

 Changes to the Service Number can be made using an administrative interface
15 for the queuing system, or through automatically generated messages as before, and the system can be configured to propagate these changes to all the other queue servers managing the different points of sale, if desired.

 Request Numbers are usually issued in ascending order for a particular domain (though some numbers may be skipped). This means that a person receiving Request
20 Numbers through one source, such as the Telephone, may be able to use that number to enter the Internet queue at a lower position.

 Note that this does not violate first-come, first-served, as the person is simply using a validly-acquired place in the queue to buy tickets through another means. However, in circumstances where this is seen as undesirable for other reasons, a
25 mechanism for periodically synchronising Request Numbers across domains or points of sale may be provided.

 Consider a situation where a (single) internet queue has issued Request Numbers 1,2,3...100, and a telephone queue has issued numbers 1,2,3...10. The telephony queue server sends a message to the internet queue server indicating that
30 it has queued 10 people. The Internet queue server reacts by skipping ten Request Numbers, so that the next new customer to arrive is issued Request Number 111.

 Subsequent messages from the Telephony queue server indicate the number of people queued since the last message sent. The Internet queue server continues to react by skipping this number of people.

35 Where multiple points of sale are involved, one point of sale server is designated the 'skipper', and the rest send 'skip' messages to that server. Usually, the

server queuing people the most rapidly should be the 'skipper', and this will usually be the Internet queue server.

It is recommended that 'skip' messages are only sent at most once every ten seconds to increase performance, however it is also possible to send 'skip' messages
5 for every new queuer if desired.

Domain Switching can benefit from integration. In cases where, for example, Internet and Telephone sales are regarded as separate domains, customers who use the internet to join the queue and the telephone to buy their tickets (or vice versa) are known as Domain Switchers.

10 It is possible to track the original domain that gave the Request Number that is being used to buy tickets across queue servers, so that when tickets are sold to Domain Switchers, the correct sales tracking information is updated, and people from the switched-to domain are not unfairly excluded from buying tickets.

In certain embodiments, the present invention can facilitate purchases of
15 groups of tickets - where this is deemed appropriate. It is often the case that groups of people all wish to attend the same event. In heavily over-subscribed sales, it may be difficult for the whole group to receive tickets. In a conventional system, the probability of the entire group being able to buy tickets is given by p^n , where n is the number of people in the group, and p is the total number of tickets available divided by
20 the total number of people who want them. For instance, in a situation where there are four times as many people who wish to go as there are tickets, the chances that a group of ten people will all receive tickets are roughly a million to one!

This may lead to problems, as the members of the group that have received tickets later decide that they do not wish to go because other members of the group
25 have been excluded. In addition to the customer dissatisfaction, this also encourages these members to sell their tickets (become touts).

If, however, the queuing system of the invention is used, and the group is organised so that they all attempt to join the queue substantially simultaneously, this may increase the chances of the group receiving tickets, however this cannot be
30 guaranteed.

To combat this problem, a subsystem of the queueing system has been developed.

To use this system, groups of people must register with the queuing system (through any point of sale) before the queue is opened (usually when tickets go on
35 sale).

The first member of the group to register is given, or chooses, a unique Group

Code, which that person shares with the rest of the group.

Subsequent members of the group then enter this Group Code into the system when they register.

5 Personal information, such as the last four numbers of each member's credit card, is also taken when the members of the group register, to prevent places in the group being traded. The maximum size of the group may be limited if desired.

When the queue opens, all members of the group should try to join it as soon as possible in order for the group to have the best chance of receiving tickets.

10 The first member of the group to join the queue is given a Request Number as before.

Subsequent members of the group joining the queue are given this same Request Number.

Once this Request Number becomes eligible for service, all members of the group should complete their purchase before tickets sell out as before.

15 If the internet is used to join the queue, it is possible for the system to recognize members of the group as they join the queue through the use of cookies or other storage technology (such as bookmarks/favourites).

20 If other points of sale are used, members may need to re-enter their group code in order to be recognized. Optionally, members may be prompted to re-enter their personal information in addition to the group code in order to discourage fraudulent use.

Note that this may violate the strictest definition of first-come, first-served, and large groups may be favoured over smaller ones. However, in physical queues, it is usual for people to save places for their friends, so this may be seen as acceptable.

25 The preceding discussion has in the main been concerned with high-load sales applications of the system of the invention. The system can also be applied to error based queuing. While most online businesses do not sell tickets, all online businesses use similar technologies to facilitate the purchase of product (goods or services). The queuing solution presented can also be used to prevent loss of business when these technologies fail.

30 Specifically, most online businesses use a web server to present web pages, a database to store product and order information, a payment gateway to process credit card details, and (usually) an application server to integrate these subsystems.

35 Problems occur when any or all of these subsystems fail. For instance, if a payment gateway is unavailable, perhaps due to a server restart, or hardware reboot, or network failure, or some other reason, people find themselves unable to make

purchases, even though the rest of the system may be working correctly. Similarly, if the database server is down, then customers will also not be able to place orders.

In these circumstances, frustrated customers will often buy product from other suppliers, and are unlikely to return to the web site they perceive as faulty. So, even
5 a routine maintenance operation, which may take down a component of the system for only a minute or two, may cause permanent loss.

The system of the invention addresses these issues. Firstly, when a component of an eCommerce site is failing, users are normally presented with some kind of error message. The specific message shown will depend on the component
10 failing and the rest of the architecture of the eCommerce site, however site administrators do usually have control over the message shown.

In short, site administrators use the queue system to replace the error message (which would usually be generated for the customer) with instructions to send the customer to the queue server. The queue system responds by placing the customer
15 in a queue. Any subsequent customers are held in this queue until the problem is fixed, at which point they are returned to the Target Service Host at a rate it can handle.

The mechanism for this is as follows.

First, when the customer is sent to the queue server (for instance, through
20 JavaScript, an HTTP redirection header or some other means), the redirection contains the following information:

- 1) Information regarding the nature of the original error (usually an error code),
- 2) A target URL to send customers to once the error is fixed,
- 3) A test URL used to determine whether the problem has been fixed,
- 25 4) A delay time,
- 5) An email address and/or telephone number for the site administrators, and/or
- 6) Any additional information required by the Site Administrators.

Any or all of these may be preconfigured to default values by the site administrators and then omitted from the redirection email. All this information is stored
30 using the processes described above.

On receiving the first customer from a failing site, the queue server sends an email and/or SMS text to the site administrators indicating that a problem has occurred that needs fixing. An error code may be included in this message.

Next, the queue server issues a Request Number to the customer and shows
35 a queue page as before. A timer is started. Subsequent customers arriving because of the error are issued further Request Numbers and join the queue as before. The

queue page is configurable, so that Site Administrators have control over any messages shown.

Once the timer has reached the delay time, the next customer to reach the queue servers (either as a new arrival, or as a customer who is already queued whose queue page is automatically refreshing) is labelled by the system as a Tester, through the use of a persistent cookie or other means. The Tester is sent to the Test URL, which is a web page on the Target Service Host (a webpage that became unavailable due to an error or failure of a component of the e-commerce system). The information stored when the Tester joined the queue is also sent along with this request to the Target Service Host. Another timer is started on the queue server.

The Target Service Host uses the sent information to test whether the failing component is now working. The Target Service Host may do this by, for instance, reconstructing the original request that resulted in error, and reprocessing it.

If the error occurs again, the customer is immediately returned to the queue server.

If the error does not recur, the Target Service Host indicates that the problem has been fixed either:

- by sending a message to the queue server directly, or
- by causing the customer's web browser to request a particular URL from the queue server that indicates that the problem has been fixed, or
- by not returning the Tester to the queue server, which detects the absence of the Tester after a certain amount of time.

If the error is still occurring, the queue server removes the information that demarcates the Tester from the rest of the customers, and continues to hold all customers for another fixed period of time, and then repeats the process.

Once the error has been fixed, however, the tester is forwarded to the Target URL either by the Target Service Host, or (optionally) by the queue server, and the rest of the customers in the queue are also sent through to the Target URL as before, until there are none left in the queue, and an email or SMS text is sent to the Site Administrators indicating that the problem has been fixed.

If the error reoccurs during this process, then the customer experiencing the error is sent to the queue server as before, which responds by holding all queued customers again until the error is fixed, using the testing mechanism described above.

Optionally, Site Administrators may also specify to the queue servers that the problem has been fixed manually, bypassing the automatic testing.

Optionally, the queue servers may also be configured to send requests to the

Target Service Host to perform testing directly, rather than by sending customers to the Test URL as described above.

5 The queuing system of the invention may therefore be deployed to address problems in both internet and telephonic queue management, where high load, dealing with multiple points of sale, and/or system failures cause inefficiency and/or failure, thereby providing reduction of risk and costs to Site Administrators, and a greatly improved experience for customers. In addition, the queuing system addresses problems arising from ticket touts, cancellations, and group ticketing. The invention is however not limited to applications in e-commerce and telephony, it encompasses the queuing of requests for services and/or products in any communications network.

0 Furthermore, although many of the specific examples given above relate to queuing systems for purchase of tickets, it will be apparent that the invention is not limited to these specific examples. The invention may equally be used in the management of queues for the purchase of products and/or services. The invention also applies to the queuing of users where access to a service is prevented by failure or malfunction rather than demand that outstrips capacity.

5 The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement of any form of suggestion that such prior art forms part of the common general knowledge.

10 It will be understood that the term "comprise" and any of its derivatives (eg. comprises, comprising) as used in this specification is to be taken to be inclusive of features to which it refers, and is not meant to exclude the presence of any additional features unless otherwise stated or implied.

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for managing requests for service to a service host over a communications network, the method comprising the steps of:

5 receiving from a customer terminal via the communications network a request for service to be supplied by the service host;

allocating a queue identifier to the request for service;

receiving from the customer terminal a subsequent request for service to be supplied by the service host;

0 retrieving the queue identifier for the request for service;

performing a comparison between the queue identifier and queue status information; and

passing to the service host the request for service in accordance with the result of the comparison,

5 wherein the steps of the method are performed at a queue server and the step of allocating the queue identifier to the request for service is performed independently of whether the request for service is currently eligible for service by the service host.

2. A method according to claim 1, further comprising the step of performing an initial comparison between the queue identifier and queue status information after the queue identifier has
0 been allocated to the request for service.

3. A method according to claim 2, further comprising the step of notifying the customer terminal that the request has been enqueued after performing the initial comparison.

25 4. A method according to claim 3, further comprising the step of providing the customer terminal with access to the queue identifier.

5. A method according to any preceding claim, wherein the queue identifier is a number and is allocated on a sequential basis.

30 6. A method according to any one of claims 1 to 5, wherein the communications network is a telephone network.

35 7. A method according to claim 6, wherein the queue identifier is stored in a database and is associated with a telephone number of the customer terminal.

8. A method according to claim 7, wherein the queue identifier is retrieved from the database following the subsequent request for service.

9. A method according to any one of claims 6 to 8, wherein the step of passing the request for service to the service host comprises connecting a call from the customer terminal to the service host.

10. A method according to any one of claims 1 to 5, wherein the communications network is the Internet.

11. A method according to claim 10, wherein the queue identifier is retrieved from the subsequent request for service.

12. A method according to claim 10 or claim 11, wherein the step of passing the request for service to the service host comprises providing the customer terminal access to a target page on the service host.

13. A method according to claim 4 or any one of claims 5 to 12 when dependent on claim 4, further comprising the step of providing the customer terminal with access to additional information relating to the queue.

14. A method according to claim 13, wherein the additional information includes a service number, the service number being an indication of the queue identifier currently being served.

15. A method according to claim 14, wherein the service number is automatically incremented at a constant rate.

16. A method according to claim 14, wherein the service number is automatically incremented, the rate of increment being varied in dependence on the number of requests for service passed to the service host.

17. A method according to claim 14, wherein the service number is incremented in response to messages from the service host.

18. A method according to claim 4 or any one of claims 5 to 17 when dependent on claim 4, wherein access to information specific to the customer terminal is provided with access to the queue

identifier.

19. A method according to any preceding claim, further including the step of checking at the service host that the request for service has been passed from the queue server.

20. A system for managing requests for service from a customer terminal to a service host over a communications network, the system comprising a server for receiving from the customer terminal via the communications network a request for service to be supplied by the service host, the server including:

means for receiving a request for service from the customer terminal;

means to allocate a queue identifier to the request for service;

means for generating queue status information;

means for retrieving the queue identifier for the request for service following a subsequent request for service from the customer terminal;

means to perform a comparison between the queue status information and the queue identifier; and

means for passing the request for service to the service host in dependence on the result of the comparison,

wherein the server is a queue server and the means for allocating the queue identifier to the request for service is adapted to allocate the queue identifier independently of whether the request for service is currently eligible for service by the service host.

1/5




Fig 1

2/5




Fig 2




Fig 3

3/5





Fig 4



5/5




Fig 6