




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁷ : H04L 29/06, G06F 9/46</p>	<p>A1</p>	<p>(11) International Publication Number: WO 00/08823 (43) International Publication Date: 17 February 2000 (17.02.00)</p>
<p>(21) International Application Number: PCT/US99/02154 (22) International Filing Date: 1 February 1999 (01.02.99) (30) Priority Data: 60/095,652 7 August 1998 (07.08.98) US (71) Applicant: E2 SOFTWARE CORPORATION [US/US]; Suite 102, 860 Avenue F, Plano, TX 75074 (US). (72) Inventors: KEENE, John, D.; 4504 Torino Place, Plano, TX 75093 (US). FARRIS, Jeffrey, L.; 4301 Lakeside Drive, Dallas, TX 75219 (US). (74) Agent: KINCART, Joseph, P.; Fish & Richardson P.C., Suite 2800, 45 Rockefeller Plaza, New York, NY 10111 (US).</p>		<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report.</i></p>

(54) Title: LOAD BALANCE AND FAULT TOLERANCE IN A NETWORK SYSTEM



(57) Abstract

In the present invention a method and apparatus is provided for load balancing and fault tolerance amongst a plurality of servers on a computer network, such as the Internet or a private network. A track of work is claimed by a server. Thereafter, the server processes work pertaining to that track and also allocates blocks of work to other servers on the network who have advertised for work. In the event a server ceases to communicate with the other servers on the network, all uncompleted work assigned to the server ceasing to communicate is reallocated to another server. The group of servers comprises a peer group, and a peer group elects a master. The presence of the master indicates that the peer group is functioning properly.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

LOAD BALANCE AND FAULT TOLERANCE
IN A NETWORK SYSTEM

5 Cross Reference To Related Applications

This application claims priority to U.S. Provisional Application No. 60/095,652, filed August 7, 1998.

Background of the Invention

10 The present invention relates to load balancing and fault tolerance amongst computer servers functioning to track Internet/Intranet transactions. In particular the present invention relates to a system of load balancing and fault tolerance utilizing a lightweight
15 algorithm and continually cycling processes to reduce exchange of server state information.

A mechanism is described to achieve both load balancing and fault tolerance. The backup systems provide load balancing services while active. When a
20 system fails, the remaining available systems take over the failed system's load. A master system determines which participating system owns a decision track. Ownership of a decision track indicates responsibility for executing a contact gathering process and an event
25 evaluation process. Step evaluation processes are distributed among available system within the same peer group.

By way of example access to distributed networks such as the internet has increased greatly in recent
30 years and challenged commerce to use the internet advantageously. Thousands of internet and intranet, hereinafter Inet, sites have been added to networks. A

- 2 -

great expenditure of time and effort has been invested in creating a myriad of resources available to Inet browsers. As a means to benefit from the Inet forum, it is useful to have tools to interact with those browsing the Inet such as being able to track parties contacting a particular Inet-site. It is important that these tools be reliable and responsive as an Inet contact may be the first and possibly only type of contact made with an individual.

10 The creation of virtual worlds online has further increased the importance of reliability and responsiveness. Purveyors of the Inet desire interactions that further emulate a real life commercial experience. Virtual storefront owners, corporate home pages, online catalogue vendors, and a myriad of other Inet-site owners, find it useful to be able to emulate the real life experience. As the complexity of an Inet interaction increases, the expectations of an individual making contact via the Inet also increases. Contact
15
20 requires a fast reliable response.

 A traditional method of increasing transaction speed is to increase the speed of processor units running the application. Processors have limits however to the maximum throughput available. Increasing demands cannot
25 always be met by a faster processing box. Another method of increasing transaction speed is through shared processing amongst a plurality of processing units. Typically, however, this has involved very complicated hardware and software solutions requiring a sizable
30 investment in man hours and expense. Often these types of solutions are not warranted for a dedicated Inet application.

 To be effective, a system needs to be right sized to the task at hand. Consequently, there remains a need

for a simple, cost effective means of sharing a processing load and also providing fault tolerance.

Summary of the Invention

Accordingly, this invention provides a method of load balancing and fault tolerance amongst a plurality of servers on a computer network, such as the internet or a private network. In a preferred embodiment of the invention, a programmed computer server divides processing work into tracks of work that may be referred to as decision tracks. Each decision track comprises a series of conditions that are to be tested by records of a database. As conditions of the decision track are tested and met, an appropriate action is taken in response to the condition met. In addition, actions can be sequenced so as to achieve a desired result, as illustrated in Table 1 below.

Decision Track		
Condition 1	If Condition is Met Then	Action 1
Condition 2	If Condition is Met Then	Action 2
Condition 3	If Condition is Met Then	Action 3
Condition 4	If Condition is Met Then	Action 4
Condition 5	If Condition is Met Then	Action 5

Decision tracks are constructed so as to be able to be claimed by a computer server. A plurality of computer servers is arranged into a peer group networked together. The network enables the computer servers to communicate with each other. Servers coordinate within the peer group to claim individual decision tracks. Thereafter, the server owning a decision track processes initial work pertaining to that track and also allocates

- 4 -

blocks of work from that decision track to other servers on the network who have advertised for work. In the event a server ceases to communicate with the other servers in the peer group, all uncompleted work assigned
5 to a non-communicating server is reallocated to another server.

Peer groups also elect a master server. The presence of a master server signifies that a group is functioning and able to handle common tasks. In addition
10 to normal peer group server functions, a master server handles tasks pertaining to the peer group as a whole, such as e-mail. The master is typically elected by a simple device such as the lowest machine number of each server.

15 For the purposes of this disclosure an Internet can refer to, for example, a network comprising computers exceeding the boundaries of a private network. An Intranet can refer to, for example, computers within a private network. An Inet can refer to an Internet and/or
20 an Intranet adhering to an internet protocol or similar protocol. An Inet-site is, for example, a site available on either an Internet or an Intranet. A network, for example, can have a computer acting as a server and a computer acting as a client. A contact can, for example,
25 be an access to an electronic interface such as a web site, or other contents of a stored memory such as a hard drive or dynamic random access memory of a server. A client can be a person, a node operator, or broadly, a machine or electronic device making such contact, or
30 causing a node of a network to make such a contact. Real time is meant to be read broadly to signify on a basis timely to or in relation to an individual event.

Other advantages and features of the present invention will become apparent from the following
35 description, including the drawings and the claims.

- 5 -

Brief Description of the Drawing

FIG. 1 illustrates a typical configuration supporting this invention.

FIG. 2 illustrates the query process of a decision track.

FIG. 3 illustrates a load balancing sequence.

Description of the Preferred Embodiments

According to the present invention, an apparatus method and system are described for load balancing and fault tolerance comprising a plurality of computer servers 110 networked together into a peer group 120 and also networked to a database server 130. The network provides a means of communication between servers. Work is divided into tracks 135 and distributed amongst the peer group servers according to the availability of each server to accommodate additional work. Utilizing a multitude of servers to process work effectively lessens the work required by a single server and effectively speeds the response of the system. The ability of a peer group to allocate work amongst available servers, and then reallocate work if a particular server should become unavailable, provides fault tolerance.

Servers periodically notify other peer group servers of their presence on a network by way of a well-known device such as a "hello" message or an "advertisement." Such advertisements are performed on a periodic cycle. A preferred periodic cycle is about 15 seconds. However, periodic cycles may be any length that is appropriate based on network characteristics, such as the number of nodes, the speed of communication, and the speed of the processor units. Generally, any periodic cycle between 5 seconds and 120 seconds is acceptable.

- 6 -

If an advertisement is not received from a server for a predetermined number of periodic cycles, such as for example, 4 cycles of 15 seconds each, the other peer group servers will consider the mute server unavailable.
5 Any work previously allocated to a server subsequently determined unavailable is reallocated amongst available servers.

In a preferred embodiment of the present invention, work is structured so that it may be executed
10 by decision tracks. Each decision track comprises a series of queries to be made against records of a database. If the conditions of a query are met, 210 then an appropriate action may be taken, if the conditions are not met, then a next record, or a next set of conditions
15 is queried.

During operation, a peer group computer server 110 will claim ownership of one or more decision tracks 135. After determination of ownership 310, a computer server 110 performs initial work such as for example, contact
20 gathering 320. Contact gathering comprises creation of a set of contact records 145 that are to be put on a particular step of the decision track 135. After the contact gathering is complete, blocks of work comprising steps are created and can either be distributed to other
25 computer servers 110 in the peer group 330 or performed by an owning server. A block of work may consist of, by way of example, a set of contact records ready for the next step of a decision track to be performed on them, or a list of steps to be executed on a particular record.
30 After the distribution, a server evaluates events 340 for any changes in conditions and cycles through the process again.

- 7 -

Distribution of work is effectuated by a response to advertisements or requests for work sent out by various computer servers 110 included in a peer group 120. As a server is capable of accepting additional work, it will send an advertisement to the other servers in the peer group requesting work, such as for example a step list block 350. The requesting computer server 350 the executes indicated steps 360. An owner computer server 110 who receives such an advertisement may send a block of work to the advertising computer server 110 to be processed. In this manner there is a continual load sharing of available work.

In one preferred embodiment, decision track ownership is claimed by attaching a claim counter to an advertisement broadcast by a server. A server will claim a decision track and set a counter to a predetermined interval, for example two. Each time the server broadcasts an advertisement, the counter decrements one. When the counter reaches zero the decision track is authoritatively owned by the claiming server. Other peer group servers may challenge the claim for a decision track by claiming it for themselves during the counter interval.

If two or more servers claim ownership of the same decision track, ownership election reverts to an arbitration routine. Arbitration determines ownership by a simple criterion such as the server with the least number of owned tracks. In the instance where two or more servers have an equal number of tracks, the ownership is awarded to the server with lowest machine ID.

A preferred embodiment teaches each server 110 maintaining a table 155 to store the time of the most

- 8 -

recent advertisement for each server and the decision tracks owned by each server. Each server queries the table to test if a predetermined period has elapsed without notification from any of the peer group servers.

5 If a predetermined period has elapsed without notification from a particular server, the non-communicating server is deemed to be unavailable. All decision tracks owned by unavailable servers are

10 reallocated to the remaining servers. Reallocation is accomplished in much the same manner as initial election. A server will advertise claiming ownership of a decision track of a server determined to be unavailable. If the advertisement is not challenged within a predetermined number of advertisement cycles, ownership is awarded to

15 the advertising server.

The allocation and reallocation process acts as fault tolerance. A decision track 135 will not be without an owner for more than the predetermined period. After the predetermined period has elapsed another server

20 110 takes ownership and the work of the decommissioned server commences again. Each peer group server 110 includes a copy of each decision track 135 as well as the table recording ownership of the various decision tracks. As a server 110 begins functioning as an owner, it

25 records ownership in the table, and commences to perform the work allocated to the owner of that decision track.

A database server 130 stores the contact data records 145 referenced in the various blocks of work performed by peer group servers executing decision

30 tracks. Typically, there is only one database server 130 from which all records are processed. In this manner all peer group servers have access to the same data.

- 9 -

A peer group 120 will also elect a master server 140. The advertised presence of a master server declares that network connectivity exists, that the peer group is communicating properly, and that operations may commence. Elections for a master server 140 are based on a simple criterion such as the lowest ID of the servers involved. In a periodic cycle the master will broadcast an "advertisement" or hello message, declaring its presence to other servers in the peer group. One preferred embodiment of a periodic cycle is 15 seconds. Another preferred embodiment of a periodic cycle is between 5 seconds and 120 seconds. The duration of the periodic cycle will depend on the speed of the network and the processing power of the servers.

If the presence of a master server 140 has not been detected by a peer group 120, through receipt of an advertisement from a master server 140 for a period of some number of periodic cycles, for example 4 cycles, the peer group elects a new master server. A period may be comprised of more or less cycles depending on the criticality of the timing for the work being performed and the processing power of the servers.

Decision tracks 135 and the criteria for each step of a decision track 135 can be created and manipulated via a user interface 165. In a preferred embodiment graphical representation for each step of a decision making process correlating to each step of a decision track is created. The graphical representation can facilitate accurate processing of data and ease of use. Another method for creating decisions tracks would include a written language statement defining criteria for each decision.

- 10 -

In a preferred embodiment of this invention, a software program on a computer readable medium is loaded on a plurality of servers. The software program comprises a front-end application that allows users to
5 access a variety of the features designed to load balance and provide fault tolerance. Features are grouped into different categories according to the type of users. A security scheme allows user access to a feature according to category.

10 An administrator can be responsible for secure configuration and maintenance of a decision track software. The administrator can configure databases and external access methods and defines access rights of various users. The administration is also responsible
15 for defining the synchronization relationships with other servers.

Decision tracks 135 may also define a series of actions to take based on different trigger events. Trigger events may be time-based single events, time-
20 based recurring events or external input and query result events. Queries may be directed to a database. In addition, queries against external Structured Query Language (SQL) accessible databases will operate. Conditionals control the transition of individual query
25 results to the next state in the decision track.

The methods and mechanisms described here are not limited to any particular hardware or software configuration, or to any particular communications modality, but rather they may find applicability in any
30 communications or computer network environment. In a preferred embodiment of this invention, a software program comprising computer readable code on a computer readable medium is loaded onto a plurality of servers.

- 11 -

The software program additionally comprises a front-end application that allows users to access a variety of the features designed to automate load sharing and fault tolerance.

5 The techniques described here may be implemented in hardware or software, or a combination of the two. Preferably, the techniques are implemented in computer programs executing one or more programmable computer that includes a processor, a storage medium readable by the
10 processor (including volatile and non-volatile memory and/or storage elements), and suitable input and output devices. The programmable computers may be either general-purpose computers or special-purpose, embedded systems. In either case, program code is applied to data
15 entered with or received from an input device to perform the functions described and to generate output information. The output information is applied to one or more output devices.

 Each program is preferably implemented in a high
20 level procedural or object-oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

25 Each such computer program is preferably stored on a storage medium or device (e.g., CD-ROM, hard disk, magnetic diskette, or memory chip) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage
30 medium or device is read by the computer to perform the procedures described. The system also may be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured

- 12 -

causes a computer to operate in a specific and predefined manner.

The invention described has broad application to a wide range of electronic interaction environments and a number of embodiments based upon the principles disclosed are possible.

What is claimed is:

- 13 -

1. A method of load balancing or fault tolerance in a system of computer servers comprising:
 - a) dividing a server workload into separate tracks of work;
 - 5 b) communicating an advertisement requesting from a server residing in a peer group of servers;
 - c) allocating a track of work to said server in said peer group requesting work;
 - 10 d) communicating on a periodic cycle the presence of each of the servers in said peer group to other servers in the peer group; and
 - e) reallocating a track of work previously allocated to a server that fails to
15 communicate on a periodic cycle within a predetermined time period.

2. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 1 wherein allocating a block of work further comprises:
 - 20 a) claiming ownership of decision tracks by servers wherein the server becomes the owner server of that decision track;
 - b) performing preliminary work on a decision track by the owner server; and
 - 25 c) transferring a block of work by the owner server to a server advertising for work.

- 14 -

3. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 1 or claim 2 wherein a track of work is a decision track.
4. A method of load balancing and fault tolerance in
5 a system of computer servers as recited in claim 1
wherein the network is an Inet.
5. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 1 further comprising election of a master server to
10 indicate a peer group is operational.
6. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 5 wherein election is based on the binary name of the peer group servers.
- 15 7. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 5, further comprising:
- a) monitoring the peer group for the presence of a master server; and
 - 20 b) electing a new master server if no master server is advertised as present for a predetermined period.

- 15 -

8. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 1 wherein the predetermined period is 4 periodic cycles of about 30 seconds each cycle.

5 9. A method of load balancing and fault tolerance in a system of computer servers as recited in claim 2 further comprising the step of storing the ownership of a decision track in a table on a peer group server.

10 10. An apparatus for providing fault tolerance or load balancing in a network of computer servers, the apparatus comprising:

- a) a peer group of computer servers networked together;
- b) a software program implementing the following
15 in said peer group of computer servers:
 - i) declare a peer group operational through the election of a master server from amongst said peer group of computer servers;
 - 20 ii) claim ownership of a decision track by one computer server of said peer group of computer servers creating an owner server;
 - 25 iii) enable the owner server to notify other servers comprising said peer group of the presence of the owner server on the network;

- 16 -

- c) request work by a peer group server other than the owner server;
- d) create a block of work by an owner server;
- e) distribute a block of work to a peer group server requesting work;
- f) monitor the presence of an owner server; and
- g) claim decision tracks previously owned by an owner server failing to be present on the network for a predetermined period.

10 11. The apparatus of claim 10 for providing fault tolerance and load balancing in a network of computer servers wherein the network is an Inet.

12. The apparatus of claim 10 for providing fault tolerance and load balancing in a network of computer servers wherein the election is based on the binary name of the peer servers.

13. The apparatus of claim 10 for providing fault tolerance and load balancing in a network of computer servers wherein the means for an owner server of notifying the other servers comprising a peer group of the presence of the owner server is an advertisement from the owner server to other servers comprising a peer group.

- 17 -

14. The apparatus of claim 10 for providing fault tolerance and load balancing in a network of computer servers further comprising:

- 5
- a) a means for monitoring the peer group for the presence of a master server; and
 - b) a means for electing a new master server if no master server is present for a predetermined period.

15. The apparatus of claim 14 for providing fault tolerance and load balancing in a network of computer servers wherein the means for monitoring a peer group for the presence of a master server is a periodic advertisement by a master.

16. The apparatus of claim 14 for providing fault tolerance and load balancing in a network of computer servers wherein the predetermined period is 4 periodic cycles of about 30 seconds each cycle.

17. Software, stored on a computer-readable medium, for gathering and disseminating information on a network, the software comprising instructions to cause a computer system to perform the following operations:

- 20
- a) provide a plurality of servers comprising a peer group;
 - b) provide communications between said servers;
 - 25 c) divide a workload into separate tracks;
 - d) communicate an advertisement from a server requesting work;
 - e) allocate a track of work to a server;

- 18 -

- 5
- f) communicate the presence of a server to the peer group on a periodic cycle; and
 - g) reallocate a track of work previously allocated to a non-communicating server if an advertisement stating the presence of said non-communicating server is not communicated to the peer group within a predetermined time period.

10 18. The article of manufacture of claim 17 further comprising:

- 15 a) computer readable code means for claiming ownership of decision tracks by a server wherein the server becomes the owner server of that decision track;
- b) computer readable code means for performing preliminary work on a decision track by the owner server; and
- 20 c) computer readable code means for transferring a block of work by the owner server to a server advertising for work

19. A programmed computer server for providing load balancing or fault tolerance in a peer group of servers the computer server comprising:

- 25 a) a memory having at least one region for storing computer executable program code;
- b) a processor for executing program code stored in said memory;
- c) a program code stored in said memory said program code implementing the following operations in said computer server:
 - 30 i) divide server workload into separate tracks of work;

- 19 -

- ii) communicate the presence of each server in the workgroup;
- iii) divide a server workload into separate tracks of work;
- 5 iv) communicate an advertisement from a server in said peer group requesting work;
- v) allocate a track of work to said server in said peer group requesting work;
- 10 vi) communicate the presence of each server comprising said peer group to other servers comprising said peer group on a periodic cycle; and
- 15 vii) reallocate a track of work previously allocated to a non-communicating server if an advertisement stating the presence of said non-communicating server is not received by the peer group within a predetermined time period.

20 20. The programmed computer of claim 19 wherein the tracks of work comprise decision tracks




Fig. 1

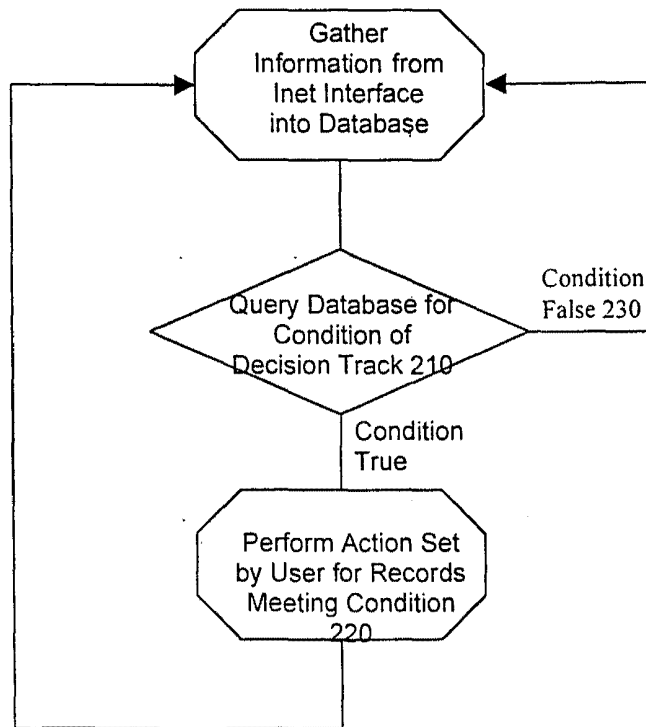


Fig. 2

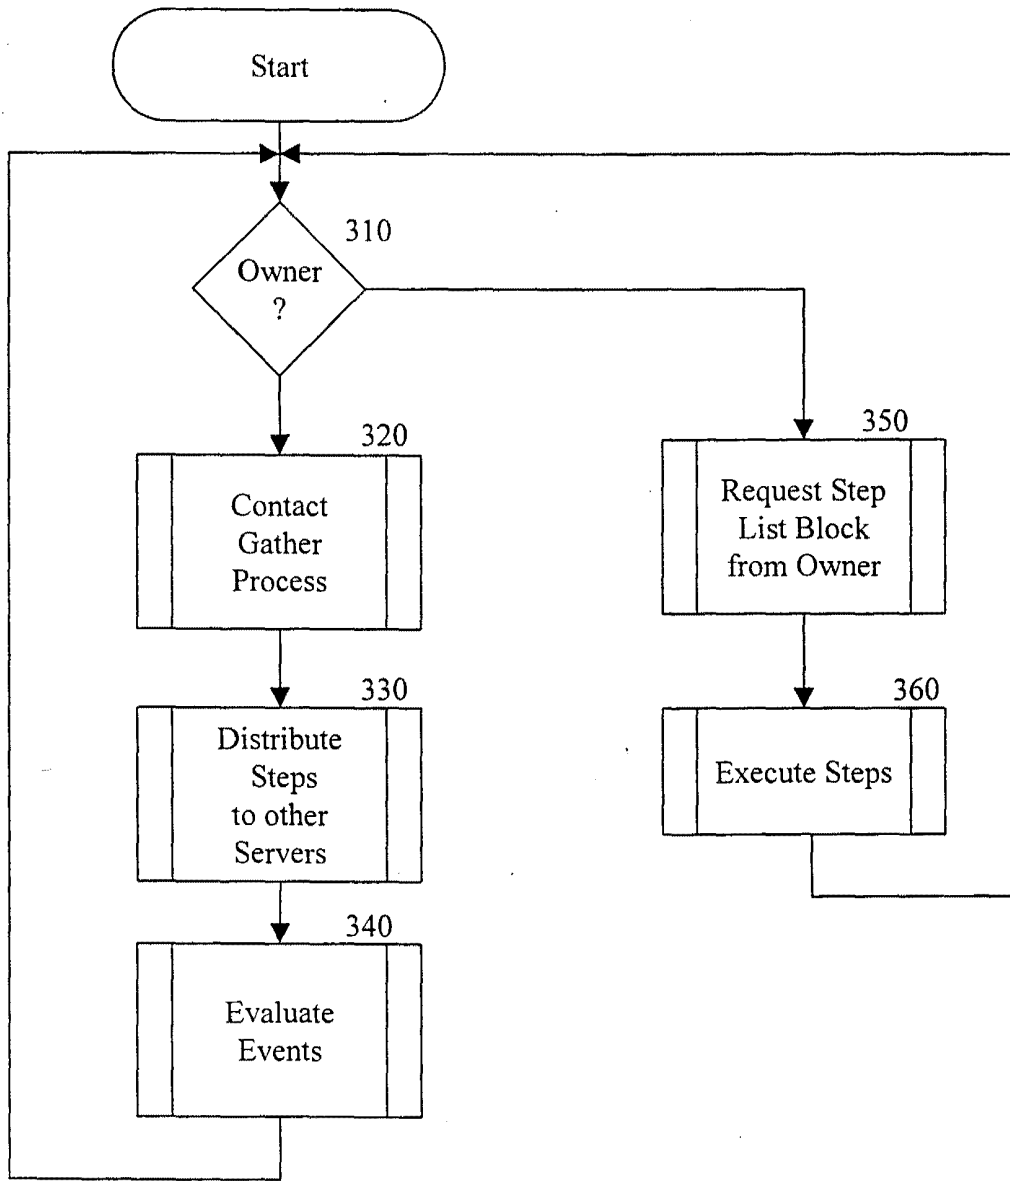


FIG. 3

INTERNATIONAL SEARCH REPORT

Inter. Appl. No.

PCT/US 99/02154

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L29/06 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	DAMANI O P ET AL: "ONE-IP: techniques for hosting a service on a cluster of machines" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 29, no. 8-13, 1 September 1997 (1997-09-01), page 1019-1027 XP004095300 ISSN: 0169-7552 abstract page 1022, right-hand column, line 32 - line 40 page 1025, left-hand column, line 40 - right-hand column, line 5 <div style="text-align: center;">--- -/--</div>	1-20



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

17 August 1999

Date of mailing of the international search report

30/08/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Adkhis, F

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 99/02154

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>YOSHIDA A: "MOWS: distributed Web and cache server in Java" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 29, no. 8-13, 1 September 1997 (1997-09-01), page 965-975 XP004095295 ISSN: 0169-7552 abstract page 966, left-hand column, line 28 - line 33 page 967, right-hand column, line 3 - line 28 page 972, left-hand column, line 10 - page 973, right-hand column, line 8 -----</p>	1-20