



US006470394B1

(12) **United States Patent**  
**Bamforth et al.**

(10) **Patent No.:** **US 6,470,394 B1**  
(45) **Date of Patent:** **Oct. 22, 2002**

(54) **APPARATUS AND METHOD FOR DATA CONVERSION AND LOAD BALANCING IN A COMPUTER NETWORK**

(75) Inventors: **John Bamforth**, Richardson; **Glenn Huber**, Eules; **Kenneth P. Thorpe, III**, Plano; **Kevin Buckley**, Grapevine; **Aravinda Shenoy**, Grapevine; **Srinivas Karra**, Grapevine; **Sudha Akshayakumar**, Plano, all of TX (US)

(73) Assignee: **The Sabre Group, Inc.**, Fort Worth, TX (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/686,998**

(22) Filed: **Oct. 12, 2000**

**Related U.S. Application Data**

(62) Division of application No. 09/031,671, filed on Feb. 27, 1998.

(51) **Int. Cl.**<sup>7</sup> ..... **G06F 13/00**

(52) **U.S. Cl.** ..... **709/241; 709/226**

(58) **Field of Search** ..... 709/246, 226, 709/105, 201, 229, 241, 224; 714/15, 4, 10

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,417,322 A	11/1983	Berry et al. ....	364/900
4,441,829 A	4/1984	Hebert, Jr. et al. ....	400/63
4,445,795 A	5/1984	Levine et al. ....	400/63
5,050,071 A	9/1991	Harris et al. ....	364/200
5,173,853 A	12/1992	Kelly et al. ....	364/419
5,202,977 A	4/1993	Pasetes, Jr. et al. ....	395/500
5,283,887 A	2/1994	Zachery .....	395/500
5,317,568 A	5/1994	Bixby et al. ....	370/85.6
5,341,477 A	8/1994	Pitkin et al. ....	395/200

5,412,566 A	5/1995	Sawa .....	364/419.14
5,434,852 A	7/1995	LaPorta et al. ....	370/58.2
5,461,611 A	10/1995	Drake, Jr. et al. ....	370/54
5,473,608 A	12/1995	Gagne et al. ....	370/85.13
5,491,693 A	2/1996	Britton et al. ....	370/85.13
5,491,800 A	2/1996	Goldsmith et al. ....	395/200.12
5,499,343 A	3/1996	Pettus .....	395/200.2
5,506,985 A	4/1996	Motoyama et al. ....	395/600
5,509,123 A	4/1996	Dobbins et al. ....	395/200.15
5,590,281 A	12/1996	Stevens .....	395/200.01
5,774,660 A	* 6/1998	Brendel et al. ....	709/201
5,774,668 A	6/1998	Choquier et al. ....	395/200.53

(List continued on next page.)

**FOREIGN PATENT DOCUMENTS**

EP	0 240 145 A	10/1987
EP	0 471867 A1	2/1992
WO	WO 97/29424 A	8/1997
WO	WO 99/44155	9/1999

**OTHER PUBLICATIONS**

Sudo, Takuma, "Infrastructure for Distributed Objects," Hitachi Review, vol. 45, No. 3, Jun. 1996, pp. 155-160.  
"Method for Universal MAC Frame Encoding," IBM Technical Disclosure Bulletin, vol. 35, No. 2, Jul. 1992, pp. 489-493.  
PCT International Search Report, PCT/US 99/04070, Dec. 15, 1999.

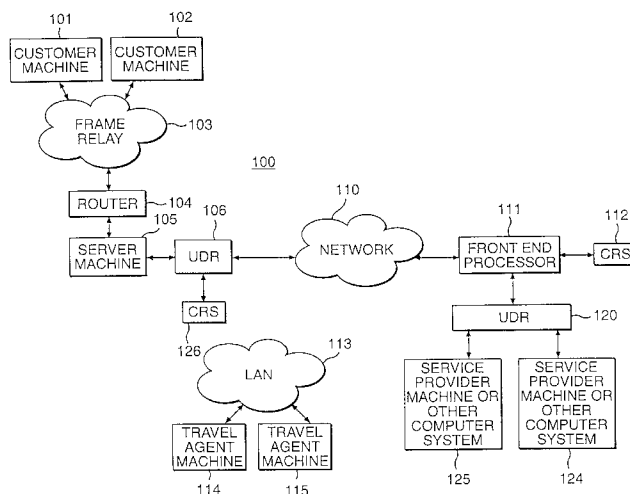
*Primary Examiner*—David Wiley

(74) *Attorney, Agent, or Firm*—Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P.

(57) **ABSTRACT**

Conversion of a data in a variable field message to a fixed format message. The variable field message has a varying length and fields of varying types of data. The corresponding fixed format message has fields of a fixed data type and length, which facilitates parsing and processing data from the variable field message. Customers for the data conversion or other processing are load balanced among multiple servers sharing load level information.

**24 Claims, 7 Drawing Sheets**



U.S. PATENT DOCUMENTS

5,828,847	A	10/1998	Gehr et al. ....	395/200.69	6,175,869	B1	*	1/2001	Ahuja et al. ....	709/226
5,832,451	A	11/1998	Flake et al. ....	705/5	6,182,139	B1	*	1/2001	Brendel .....	709/226
5,862,348	A	* 1/1999	Pedersen .....	395/200	6,185,601	B1	*	2/2001	Wolf .....	709/203
5,867,706	A	* 2/1999	Martin et al. ....	395/675	6,185,619	B1	*	2/2001	Joffe et al. ....	709/229
5,951,694	A	* 9/1999	Choquier et al. ....	714/15	6,249,801	B1	*	6/2001	Zisapel et al. ....	709/105
6,078,960	A	* 6/2000	Ballard .....	709/229	6,272,523	B1	*	8/2001	Factor .....	709/201
6,128,279	A	* 10/2000	O'Neil et al. ....	370/229	6,279,001	B1	*	8/2001	DeBettencourt .....	707/10
6,138,159	A	* 10/2000	Phaal							

\* cited by examiner

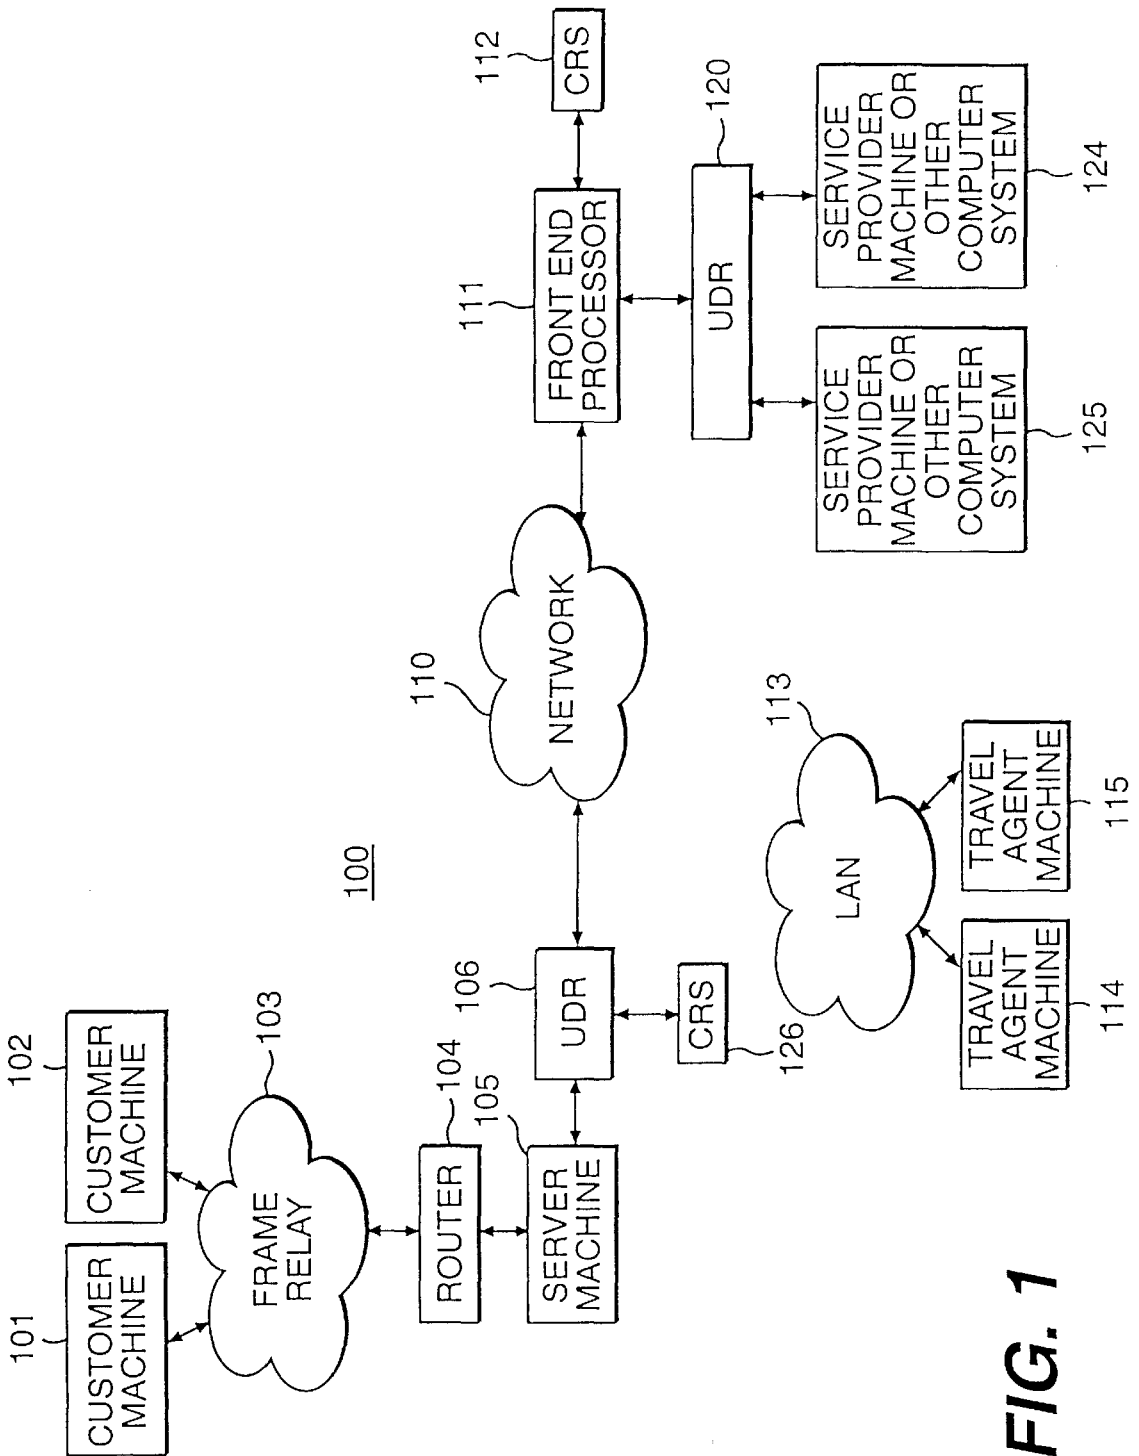


FIG. 1


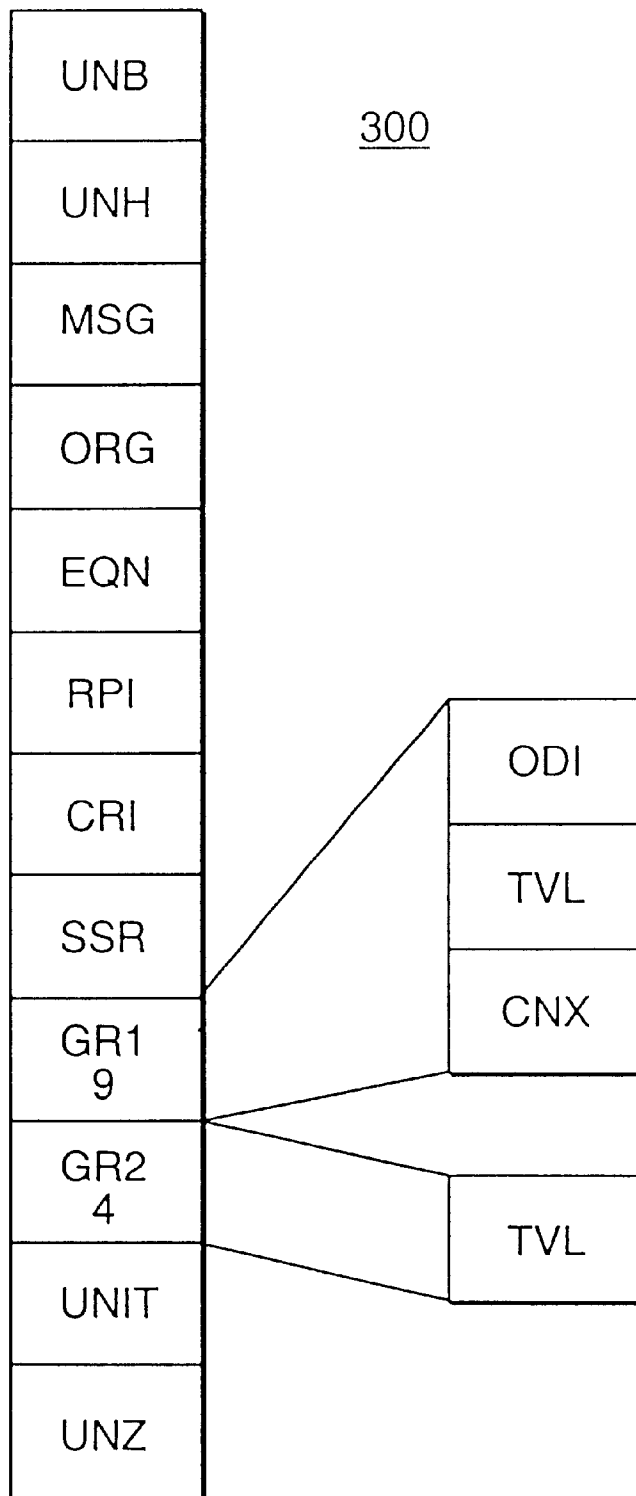



FIG. 2



**FIG. 3**



**FIG. 4**



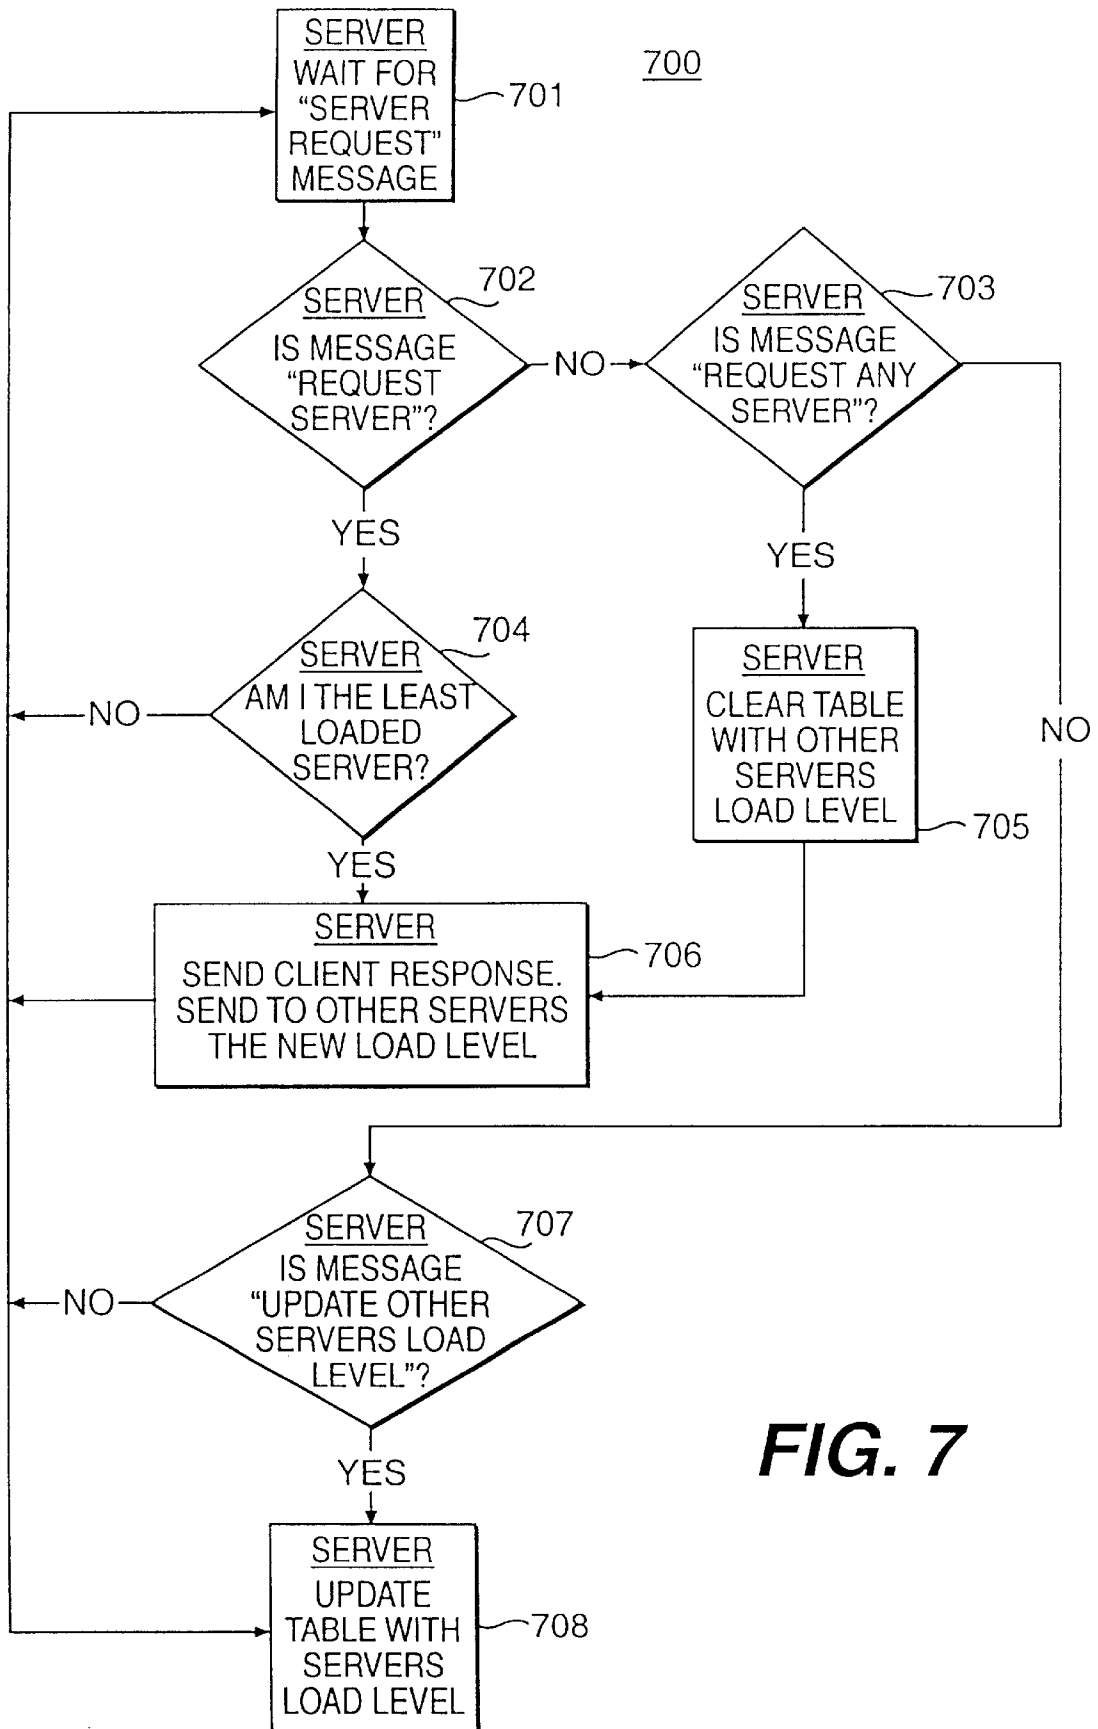


FIG. 5



**FIG. 6**



**FIG. 7**

1

## APPARATUS AND METHOD FOR DATA CONVERSION AND LOAD BALANCING IN A COMPUTER NETWORK

This is a division of application Ser. No. 09/031,671, filed Feb. 27, 1998, which is incorporated herein by reference.

### FIELD OF THE INVENTION

The present invention relates to an apparatus and method for performing conversion of data and for load balancing between machines in a computer network.

### BACKGROUND OF THE INVENTION

A computerized reservation system (CRS) traditionally has provided a communications network for travel agents or other persons to book airline reservations. Other companies may interface their computer systems with a CRS in order to make information concerning their services available via the CRS. For example, a hotel company may interface its reservation system with a CRS so that when a person books an airline reservation, he or she may also make a hotel reservation through the same network.

CRS's typically use a complex protocol for identifying and transferring data. Other companies desiring to interface their computer systems with a CRS, therefore, may find it difficult or expensive to modify their computer systems to accommodate the complex protocol of a CRS. This situation may discourage other companies from interfacing their computer systems with a CRS, which limits the available information via that network.

In addition, customers interfacing their computer systems with a CRS are typically routed through servers (or other machines) interfacing with the CRS. Generally, different servers are assigned to particular customers and service only those customers, each customer being typically routed through only one server. Therefore, if that server becomes nonfunctional, service to the customer is unavailable until the server can be brought back into operation or back on-line.

Accordingly, a need exists for data conversion and for load balancing access to a CRS or other computer network.

### SUMMARY OF THE INVENTION

An apparatus consistent with the present invention converts message format. The apparatus receives an input message in a variable field format and identifies data types for each field of the input message. The apparatus assembles an output message having multiple fields defined by a fixed format, and it maps data from the input message to specified fields in the output message based on the identified data types.

A method consistent with the present invention provides for computer-implemented message format conversion. The method receives an input message in a variable field format and identifies data types for each field of the input message. The method assembles an output message having multiple fields defined by a fixed format, and it maps data from the input message to specified fields in the output message based on the identified data types.

Another apparatus consistent with the present invention provides for load balancing of processing provided by multiplier server machines interfaced with a client machine servicing multiple customers. The apparatus assigns each of the customers to a set of the server machines and stores load

2

levels of each of the server machines. The apparatus services the customers using the server machines to which they are assigned based upon the load levels. If the server machines to which they are assigned are unavailable, the apparatus services the customers using the server machines to which they are temporarily assigned during the unavailability.

Another method consistent with the present invention load balances processing provided by multiple server machines interfaced with a client machine servicing multiple customers. The method assigns each of the customers to a set of the server machines and stores load levels of each of the server machines. The method services the customers using the server machines to which they are assigned based upon the load levels. If the server machines to which they are assigned are unavailable, the method services the customers using the server machines to which they are temporarily assigned during the unavailability.

### BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings are incorporated in and constitute a part of this specification and, together with the description, explain the advantages and principles of the invention. In the drawings,

FIG. 1 is a diagram of an exemplary computer network in which systems consistent with the present invention may be implemented including multiple CRS's;

FIG. 2 is a diagram of an exemplary apparatus for performing data conversion;

FIG. 3 is a data structure diagram representing an example of a correlation between a variable field protocol and a fixed field protocol;

FIG. 4 is a flow chart of a process for performing data conversion from a variable field protocol to a fixed field protocol;

FIG. 5 is a diagram of an exemplary apparatus for performing load balancing;

FIG. 6 is a flow chart of a process for a client machine to interface with server machines performing load balancing; and

FIG. 7 is a flow chart of a process for server machines performing load balancing to interface with a client machine.

### DETAILED DESCRIPTION

The following detailed description of the invention refers to the accompanying drawings. While the description includes exemplary embodiments, other embodiments are possible, and changes may be made to the embodiments described without departing from the spirit and scope of the invention. The following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and their equivalents.

FIG. 1 is a diagram of an exemplary network 100 including multiple CRS's. CRS's are networks permitting access to, for example, travel-related information for making reservations or obtaining such information, and CRS's may use and provide other types of information, depending upon the computer systems interfaced with a particular CRS or the information accessible by the CRS. CRS's are also referred to as computer reservation systems or central reservation system. In European countries, for example, CRS's are often referred to as global distribution systems. The term "computerized reservation system" and the abbreviation "CRS" are intended to encompass computerized reservation systems, computer reservation systems, central reservation

systems, and global distribution systems. Examples of CRS's include those known by the following trademarks and companies: SABRE; AMADEUS; WORLDSPAN; SYSTEM ONE; APOLLO; GEMINI; GALILEO; and AXESS.

Network 100 illustrates how customers or service providers may be linked together through a CRS 112 or 126. For example, customer machines 101 and 102 may represent machines located at particular companies or other entities for providing travel-related and other services for that corporation or entity. Customer machines 101 and 102 are typically interfaced through a frame relay 103 and a router 104 to a server machine 105. Router 104 provides for routing of a protocol over frame relay 104 for long distance communication. Frame relays and routers are known in the art. Server machines 105 provides necessary interaction between the ultimate customer machines and a CRS, for example, CRS 126.

Server machine 105 is typically interfaced through a universal data router (UDR) 106 to a network 110. UDR 106 may include several servers, as explained below, for performing data conversion for server 105 to communicate with a CRS, for example, CRS 126. Network 110 may represent a private network such as the Societe Internationale Telecommunications Aeronautiques (SITA) network. Network 110 interfaces UDR 106 with a front end processor 111, which provides an interface to a CRS 112. CRS's usually include a front end processor, which are known mainframe components, providing functionality for interfacing the CRS with a network. Customer machines 101 and 102 may also be interfaced with other CRS's 126 through UDR 106. Therefore, when a person at customer machine 101 or 102 desires to, for example, book a travel-related reservation or access other types of information, a communications link is established through the various elements between the customer machine and CRS 112 or 126.

In addition, network 110 may interface travel agent machines with CRS 112 or 126. In particular, network 110 may interface a local area network (LAN) 113 connected to travel agent machines 114 and 115. Travel agent machines 114 and 115, if located overseas, may also be linked into CRS 112 or 126, in which case network 110 may interface token ring LAN 113 through an international telephone or computer network (not shown). Travel agent machines and LANs are known in the art.

Other companies or service providers may also provide information available via CRS 112. They provide such information by interfacing service provider machines (or other computer systems) 124 and 125 through UDR 120 to front end processor 111. UDR 120, which may include several servers as explained below, provides data conversion to interface the computer systems of service provider machines 124 and 125 with the protocol such as CRS 112. Alternatively, service provide machines 124 and 125 may interface with UDR 106 or CRS 126.

#### Data Conversion

Data conversion is useful, for example, for converting messages in a variable field format into a fixed format, making it easier for service providers to communicate electronically with a CRS that employs the Edifact protocol, explained below, based on a variable field message format. CRS's typically operate using the Edifact protocol, which is complex. A challenge for companies operating CRS's who want to add other computer systems to the CRS involves making it cost-effective for the other computer systems to interface with data in the Edifact protocol. A data conversion

explained below, for converting from the Edifact protocol to a fixed format message provides an advantage of reducing expense and time for a service provider or other company to interface its computer system with a CRS in comparison to directly with data in the Edifact protocol. An embodiment, consistent with the present invention, converts data in the Edifact protocol to a fixed format message. The availability of information in a fixed format message, rather than having to interface directly with data in the Edifact protocol, significantly reduces the complexity of interfacing computer systems with the CRS. Interfacing other computer systems with a CRS increases the information available through the CRS and hence may increase the usefulness of the network including the CRS and the systems interfacing it.

FIG. 2 is a diagram of a data converter 201 for performing conversion of data. This conversion may be performed, for example, in UDR 120 between CRS 112 and service provider machines 124 and 125 shown in FIG. 1. As shown in FIG. 2, the conversion generally occurs between a client machine 207 and server machine 204. Client machine 207 interfaces with a plurality of terminals 208, 209, and 210. Client machine 207 may represent, for example, a computer system of a service provider, for example, service provider machine 125, providing travel-related services or access to other information, and terminals 208, 209, and 210 may represent computer terminals for users to interact with the computer system, for example, client machine. Server machine 204 interfaces with a CRS 205 and possibly another network 206. Alternatively, data converter 201 may interface directly with a CRS as shown by the dashed line. Data converter 201 includes a conversion engine 203 interfacing a function library 202 for retrieving particular software functions or scripts for performing data conversion depending on the segments or fields within a received message.

As identified above, CRS's, such as CRS 205, typically use a known protocol referred to as the Edifact (electronic data interchange for administration, commerce and transport) protocol or standard. Various versions of the Edifact protocol exist, such as those known as IOTA and UN standards, and the Edifact protocol is used generally for processing data for travel, banking, and industrial purposes. An example of the Edifact protocol is explained in the following document, which is incorporated herein by reference: Henry Schlieper, "Henry's Yellow UN/EDIFACT Book; Introduction to UN/EDIFACT Messages," 10th revised issue (October 1996). The term "Edifact" in this description is intended to cover any version of the Edifact protocol used for any purposes. With reference to FIG. 2, server machine 204 processes, for example, data formatted according to the Edifact protocol, which is complex. Client machine 207, on the other hand, may use its own protocol for the data on which it operates, particularly if client machine 207 represents the computer system of a service provider that does not support the CRS protocol. Data converter 201 provides a fixed format message to client machine 207. Thus, data converter 201 thus may provide a more simple interface for client machine 207 such that client machine 207 need not analyze a message in the complex Edifact or other protocol.

The Edifact protocol is in some respects complex because it includes variable fields. In other words, a particular type of data may appear in different fields depending upon each particular Edifact message, which is data formatted according to the Edifact protocol. There are different types of Edifact messages and examples includes, but are not limited to, the following: Availability Request Message (AVLREQ), Availability Response Message (AVLRSP), Profile Request

message (PROREQ), Profile Response Message (PRORSP), Reservation Request Message (RESREQ) and Reservation Response Message (RESRSP).

In general, a variable field message includes data formatted according to a variable field protocol. Edifact or other variable field messages may not be of fixed length and, depending on the data to be transmitted, the length of Edifact or other variable field messages may change. Thus, a machine receiving an Edifact or other variable field message cannot necessarily determine that certain data is of a particular type because of the field in which it exists. In order to understand an Edifact or other variable field message, a machine typically has to analyze each segment of data in the message to determine the type of data.

A fixed format message is a record or other data structure in which each particular field is known or predefined so that in a stream of data, for example, each field has a certain length and type of content. With data in a fixed format message, therefore, a machine need only, for example, count bytes to determine fields and data types. Thus, in a fixed format message, the data for each segment is typically at a particular position or offset in the message, which facilitates a customer's ease in processing the data and may result in a faster rate of data manipulation than if the customer were to analyze the Edifact or other variable field message.

A fixed format message data structure may be of fixed length for each type of Edifact or other variable field message. A structure for a fixed format message is typically defined depending on a particular Edifact or other variable field message and data required by a customer or a particular application. For each Edifact or other variable field message data structures, one or more fixed format messages may be defined. A fixed format message may be defined by analyzing the Edifact or other variable field message data structure for each type of message and may depend upon the size and characteristics of the messages. The description of data fields for the structure of a fixed format message may change depending on the type of data received or transmitted. It is possible to define one fixed format message for all Edifact messages. However, that fixed format message would be large, containing fields that may not be used by many Edifact messages and potentially adversely processing speed because of its size. Therefore, it may be more advantageous to define a fixed format message for a group of Edifact messages, for example, ten to fifteen Edifact messages for one fixed format message.

FIG. 3 is an example of a data structure 300 illustrating a correlation between segments of an Edifact or other variable field message and corresponding fields of a fixed format message. The letters within each box represent a known segment of an Edifact message, and each box has a predefined length in bytes. The linking of boxes illustrates how the fields are linked together and a definition of each one. The number nine in the box for the group 1 (GR1) segment means that the corresponding group of three segments (ODI, TVL, CNX) are repeated nine times. The number four in the box for the group 2 (GR2) segment means that the corresponding segment (TVL) is repeated four times. This example illustrates how variable fields of an Edifact or other variable field messages may be translated into a fixed format message so that a system reading such a message knows the type of data in each field and its length. Table 1 provides an explanation of each three-letter Edifact code in the data structure of FIG. 3.

TABLE 1

	code	meaning
5	UNB	interchange header
	UNH	message header
	MSG	message segment
	ORG	origination of request details
	EQN	number of units
	RPI	related product information
10	CRI	consumer reference information
	SSR	special requirements details
	ODI	origination and destination details
	TVL	travel product information
	CNX	connection details
	UNT	message trailer
15	UNZ	interchange trailer

FIG. 4 is a flow chart of a process 400 for performing a conversion between an Edifact or other variable field message and a fixed format message. Conversion engine 203 (FIG. 2) includes a processor for implementing the conversion process and retrieving software functions, referred to as scripts, from function library 202 for performing data conversion. For example, to convert each type of fixed format message one or more scripts may be used to perform the data conversion. An advantage of using scripts, or similar software processing, is that they may be modified and reloaded without a modification to the corresponding machine code. Conversion engine 203 may compile scripts at run time and, depending on the data received, it may select the corresponding script or other software function and assemble a fixed format message from data in a variable field message.

As shown in FIG. 4, conversion engine 203 first receives an Edifact or other variable field message, typically including protocol information, from a CRS or other computer system (step 401). It may determine a type of the message based on predefined criteria such as, for example, requirements for the Edifact protocol. Conversion engine 203 reads a segment or field in the message (step 402), and it determines or identifies the type of data in that segment or field (step 403). The determination or identification may involve using the type of message, and it may also involve downloading predefined software functions from function library 202 to analyze particular message segments or fields. Conversion engine 203 determines the location of that type of data in the corresponding fixed format message (step 404), which is typically predefined so that it knows the structure and format of that type of message. Conversion engine 203 may also remove unnecessary protocol information from the data (step 405). The protocol information is generally not required in the fixed format message, as the size and data type of the fields may be predefined or known. Conversion engine 203 assembles fixed format message and maps data from the Edifact or other variable field message to the fields in the fixed format message (step 406) by positioning the data, typically without the protocol information, in the corresponding fields for that data in the fixed format message.

Conversion engine 203 determines if the Edifact or other variable field message contains more segments or fields to process (step 407). If so, it may repeat steps 402-406 to process those segments or fields. Once the message has been translated and the fixed format message has been assembled, conversion engine 203 sends the assembled fixed format message to a client machine or other computer system or network (step 408), and the system may send it in serialized form.

The following is example of a correlation between an Edifact message and a fixed format message. This example

is provided for illustrative purposes only. Any type of fixed format message may be defined for an Edifact or other variable field message. An example of an Edifact message availability request is shown in Table 2.

TABLE 2

UNB + UN:1 AALARES:IEDI + XXXX:IEDI + 920130:1330 + SES0001 ++ AVRL1 + E'
UNH + 1 + AVLREQ:95:1:IA + 574820'
MGS + 1:29'
ORG + AA:HDQ + 31599253:A0B0 +++ 1 + US'EQN + 4:9*2:10'
CRI + 6:8*6:10'
ODI + DFW*MBJ'
TVL + 24121995:1115 + DFW*MBJ + DL +++ BCO'
CNX + MIA'
ODI + MBJ*DFW'
TVL + 12011995:1300 + MBJ*DFW + DL +++ BCO'
CNX + MIA'
UNT + 12 + 1'
UNZ + 1 + SES---1---1'

The Edifact message shown in Table 2 contains twelve segments such as, for example, UNB and UNH. Each segment may have multiple composites, each composite separated by a “+” symbol. Each element inside a composite may be separated as a “:” symbol. If a segment contains repeating elements, they may be separated by a “\*” symbol. The data inside an Edifact or other variable field message may vary depending on the type of message and the required information.

The data in Table 3 shows an example of a correspondence between a fixed format message and the first three segments (UNH, UNB, MSG) of the Edifact message shown in Table 2. Table 3 provides the data field definitions in the first column along with the offset and the size of the data. The other segments may be defined in a similar manner. All data fields may be null terminated to obtain the defined length. For example, the first element in this data structure has a data field size of “5.” The data is “UN” and the rest of the field includes three null characters (for example, “0”) to obtain a five character field. Other characters may be used to terminate fields. Accordingly, the following is a serialized version of the first five fields of fixed format message shown in Table 3: “UN00010AALARES0000IEDI0XXXX0000000”. The remaining fields may be attached in a similar manner by using null termination to obtain the appropriate field lengths.

TABLE 3

Edifact segment	fixed field	format length	data
AVRQ-UNB.CTRLAGNCYCODE	Size:	5	'UN'
AVRQ-UNB.SYNTAXVERNO	Size:	2	'1'
AVRQ-UNB.INTAPPTITLE	Size:	11	'AALARES'
AVRQ-UNB.INTADDRVERSION	Size:	5	'IEDI'
AVRQ-UNB.APPENTITYTITLE	Size:	11	'XXXX'
AVRQ-UNB.ADDRVERSION	Size:	5	'IEDI'
AVRQ-UNB.DATEGMT	Size:	7	'920130'
AVRQ-UNB.TMEGMT	Size:	5	'1330'
AVRQ-UNB.INTCTLREF	Size:	15	'SES00010001'
AVRQ-UNB.APPASSCID	Size:	15	'SES00010001'
AVRQ-UNB.FSEID	Size:	15	'AVRL1'
AVRQ-UNB.ASSOCCODE	Size:	2	'E'
AVRQ-UNH.MSGKEFNO	Size:	14	'1'
AVRQ-UNH.MSGTYPE	Size:	7	'AVLREQ'
AVRQ-UNIA.MSGVERNO	Size:	3	'95'
AVRQ-UNH.MSGRLSNO	Size:	2	'1'
AVRQ-UNH.CNTLAGENCY	Size:	3	'IA'
AVRQ-UNH.COMACESREF	Size:	18	'574820'

TABLE 3-continued

Edifact segment	fixed field	format length	data
AVRQ-MSG.BUSFUNCCODE	Size:	4	'1'
AVRQ-MSG.MESSFUNCCODE	Size:	4	'29'
AVRQ-MSG.RESPCODECNT	Size:	4	'0'
AVRQ-MSG-RESPCODE.RESPCTYPE	Size:	3	'0'
AVRQ-MSG-RESPCODE.RESPCTYPE	Size:	3	'29'
AVRQ-MSG-RESPCODE.RESPCTYPE	Size:	3	'0'
AVRQ-MSG-RESPCODE.RESPCTYPE	Size:	3	'0'
AVRQ-MSG-RESPCODE.RESPCTYPE	Size:	3	'0'

Clustering Servers for Load Balancing

Load balancing involves assigning customers to one or more servers for performing processing for the customers. Balancing the customer load among the servers is important, for example, to maintain service to the customers and avoid downtime in which service is unavailable.

FIG. 5 is a block diagram illustrating how load balancing may be accomplished in a CRS or computer network. Machine or network 506 interfaces with client machine 501 through, in this example, four servers. These servers include server 502 servicing customer A, server 503 servicing customers A and B, server 504 servicing customer B and C, and server 505 servicing customers C and D. Each customer is thus assigned to one or more servers interfacing client machine 501 with machine or network 506. In addition, each server is typically assigned to a particular port and constitutes the address of that server, as shown by the exemplary port numbers. For example, server 502 is assigned to port 1000 (p1000), and the other servers are assigned in this example to sequential port numbers. Each server typically includes an element, such as a software table or other data structure, for storing load levels for each of the servers, explained below. As shown in FIG. 5, servers 502, 503, 504, and 505 contain, respectively, load level tables 507, 508, 509, and 510. Each server may have a table or other data structure indicating its own load level and the load levels of all other servers. Load levels may indicate, for example, how many customers a particular server machine is currently servicing. Using the load level information provides for additional load level balancing of customer processing by, for example, routing customers through the server with the lowest load level.

Therefore, load balancing may in general be accomplished using the assignment of customers to servers machines and the load level information. In particular, as mentioned above, customers may be assigned for service to a set of the servers, and each server may store load levels of each of the servers. Customers may be serviced using the servers to which they are assigned based upon the load levels. In particular, among the servers to which they assigned, they may be serviced by the server with the lowest load level. If the servers to which they are assigned are unavailable, the customers may be serviced using the servers to which they are temporarily assigned during the unavailability.

Servers 502–505 may represent, for example, the UDRs explained above. With reference to FIG. 1, client machine 501 may represent server machine 105 or service provider machines 124 and 125, and customers may represent customer machines 101 and 102, or service provider machines 124 and 125. Machine or networks 506 may represent network 110, front end processor 111 in combination with CRS 112, or CRS 126. Thus, servers 502–505 may provide, for example, data conversion necessary for customer machine 101 and 102, or service provider machines 124 and

125, to interface with CRS's 126 and 112. The load bearing achieved by servers operating consistent with the present invention may be used in any applicable computer network and for any applicable processing.

FIG. 6 is a flow chart of a process 600 for client machine 501 to interface with servers 502-504 performing load balancing. Client machine 501 contains a processor and memory for performing the load balancing function. As shown in FIG. 6, the client machine sends a "request server" message, indicating that one of the customer machines needs service by one of the servers to which the customer is assigned (step 601). The client machine may broadcast the message to the ports of servers to which the customer is assigned. The client machine determines if it received a response (step 602). If it received a response, it connects to the server specified in the response (step 603). Otherwise, if it received no response, the client machine sends a "request any server message" (step 604). The client machine broadcasts this message to the ports for all servers.

FIG. 7 is a flow chart of a process 700 for servers 502-504 to interface with client machine 501. Each of the servers contains a processor and memory for performing this process. As shown in FIG. 7, each server is in a wait state waiting for a "server request" message from a client machine through an assigned port (step 701). When the server receives a message, it determines if the message is a "request server" message, meaning that a customer assigned to that particular server requested service (step 702). If so, the server determines if it is the least loaded server, which may be accomplished by evaluating its table of server load levels (step 704). The load levels may be stored in tables 507-510 for respective servers 502-505 (see FIG. 5). Load levels may constitute, for example, numerical values indicating how many customers a server is currently servicing.

If the server that received the "server request message, is the least loaded server, it sends a response to the client machine and sends to the other servers its new load level (step 706). The response indicates to the client machine that this particular server is available for servicing the customer requesting service. If the server is no the least loaded server, then it waits for another "server request" message (step 701).

If the received message was not a "request server" message, the server determines if the message is a "request any server" message (step 703). If so, the server clears its table containing the other servers load levels (step 705), and it sends a response to the client machine and sends to the other servers its new load level (step 706).

If the received message is neither of those messages, the server determines if the message is an "update other servers load level" message (step 707). If so, the server updates its table with the other servers new load levels (step 708) and waits for another message (step 708).

Additional servers may join the servers already servicing a group of customers. When a server joins the others, it may be predefined to service a group of customers and may broadcast a message to the other servers identifying itself and its load level. It then may perform the processing described above with respect to FIG. 7.

While the present invention has been described in connection with a preferred embodiment, many modifications will be readily apparent to those skilled in the art, and this application is intended to cover any adaptations or variations thereof. For example, various other components for the elements shown in FIG. 1, different types of variable field messages and fixed format messages for data conversion, and more or fewer servers for load balancing may be used without departing from the scope of the invention. This invention should be limited only by the claims and equivalents thereof.

What is claimed is:

1. An apparatus for load balancing of processing provided by multiple server machines directly interfaced with a client machine, comprising:

means for storing at the server machines load levels of each of the server machines;

means for assigning the client machine to a first server machine among a predetermined set of server machines based upon the load levels; and

means for temporarily reassigned the client machine to a second server machine not among the predetermined set of server machines based upon the load levels if each of the predetermined set of server machines is unavailable.

2. The apparatus of claim 1, further comprising:

servicing means for receiving a message from the client machine requesting service at at least one of the first server machine to which the client machine is assigned and the second server machine to which the client machine is temporarily assigned.

3. The apparatus of claim 2, wherein the servicing means includes

means for determining if a particular one of the server machines among the predetermined set of server machines is least loaded among the server machines, and

means for sending a response if the particular server machine is the least loaded server machine.

4. The apparatus of claim 2, wherein the servicing means includes

means for sending a new load level for a particular one of the server machines to the other server machines.

5. The apparatus of claim 4 wherein the storing means further comprises means for updating the load levels at each of the server machines with the new load level.

6. The apparatus of claim 2, wherein the servicing means includes:

means for receiving an input message in a variable field format;

means for identifying data types for each field of the input message;

means for assembling an output message having multiple fields defined by a fixed format; and

means for mapping data from the input message to specified fields in the output message based on the identified data types.

7. A system for load balancing of processing provided by multiple server machines interfaced with a client machine servicing multiple customers, comprising:

a plurality of server machines interfaced directly with the client machine; and

an apparatus servicing the customers by the server machines, the apparatus comprising:

means for assigning the client machine to a first server machine among a predetermined set of server machines based upon the load levels; and

means for temporarily reassigning the client machine to a second server machine not among the predetermined set of server machines based upon the load levels if each of the predetermined set of server machines is unavailable.

8. The system of claim 7, wherein the apparatus servicing the customers further comprises:

servicing means for receiving a message from the client machine requesting service at at least one of the first

11

server machine to which the client machine is assigned with the second server machine to which the client machine is temporarily assigned.

9. The system of claim 8, wherein the servicing means includes

means for determining if a particular one of the server machines among the predetermined set of server machines is least loaded among the server machines, and

means for sending a response if the particular server machine is the least loaded server machine.

10. The system of claim 8, wherein the servicing means includes

means for sending a new load level for a particular one of the server machines to the other server machines.

11. The system of claim 10 wherein the storing means further comprises means for updating the load levels at each of the server machines with the new load level.

12. The system of claim 8, wherein the servicing means includes:

means for receiving an input message in a variable field format;

means for identifying data types of each field of the input message;

means for assembling an output message having multiple fields defined by a fixed format; and

means for mapping data from the input message to specified fields in the output message based on the identified data types.

13. A method for load balancing of processing provided by multiple server machines directly interfaced with a client machine servicing multiple customers, comprising the steps of:

assigning the client machine to a first server machine among a predetermined set of server machines based upon the load levels; and

temporarily reassigning the client machine to a second server machine not among the predetermined set of server machines based upon the load levels if each of the predetermined set of server machines is unavailable.

14. The method of claim 13, further comprising:

servicing the customers by receiving a message from the client machine requesting service at at least one of the first server machine to which the client machine is assigned and the second server machine to which the client machine is temporarily assigned.

15. The method of claim 14, wherein the servicing step includes the steps of

determining if a particular one of the server machines among the predetermined set of server machines is least loaded among the server machines, and

sending a response if the particular server machine is the least loaded server machine.

16. The method of claim 14, wherein the servicing step includes

the step of sending a new load level for a particular one of the server machines to the other server machines.

17. The method of claim 16 wherein the storing step further comprises the step of updating the load levels at each of the server machines with the new load level.

18. The method of claim 14, wherein the servicing step includes the steps of:

receiving an input message in a variable field format;

identifying data types for each field of the input message;

assembling an output message having multiple fields defined by a fixed format; and

12

mapping data from the input message to specified fields in the output message based on the identified data types.

19. A method for load balancing of processing provided by a plurality of server machines for a client machine, the plurality of server machines being divided into a predetermined group of one or more server machines assigned to the client machine and a second group of one or more server machines not assigned to the client machine, comprising:

establishing a load level for each of the predetermined group of server machines;

broadcasting the load level from each of the predetermined group of server machines to all others of the predetermined group of server machines so that each server machine in the predetermined group may store a load level for all server machines of the predetermined group of server machines;

broadcasting a request for service message from the client machine to the predetermined group of server machines;

providing service from the server machine with the lowest load level from the predetermined group of server machines to the client machine;

establishing a load level for each of the second group of server machines if no service is provided by the predetermined group of server machines;

broadcasting the load level from each of the second group of server machines to all others of the second group of server machines so that each server machine in the second group may store a load level for all server machines of the second group of server machines if no service is provided by the predetermined group of server machines;

broadcasting a request for service message from the client machine to the second group of server machines if no service is provided by the predetermined group of server machines; and

providing service from the server machine with the lowest load level from the second group of server machines to the client machine if no service is provided by the predetermined group of server machines.

20. The method of claim 19 further comprising broadcasting an updated load level from the server machine providing service to all others among a group of server machines so that each server machine may update its stored load level for the server machine providing service.

21. The method of claim 20 further comprising broadcasting an updated load level from a server machine which has finished providing service to all others among a group of server machines so that each server machine may update its stored load level for the server machine which has finished providing service.

22. The apparatus of claim 1 further comprising means for servicing the client machine using server machines to which the client machine is not assigned if the set of server machines to which the client machine is assigned are unavailable for service.

23. The system of claim 7 wherein the apparatus for servicing the customers further comprises means for servicing the customers using server machines to which the customers are not assigned if the set of server machines to which the customers are assigned are unavailable for service.

24. The method of claim 13 further comprising servicing the customers using server machines to which the customers are not assigned if the set of server machines to which the customers are assigned are unavailable for service.

\* \* \* \* \*

UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 6,470,394 B1  
DATED : October 22, 2002  
INVENTOR(S) : John Bamforth et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 10,

Line 10, "reassigned" should read -- reassigning --.

Column 11,

Line 2, "with" should read -- and --.

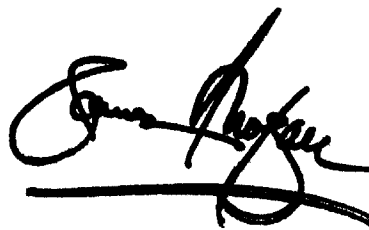
Line 22, "of each" should read -- for each --.

Column 12,

Line 4, "serve" should read -- server --.

Signed and Sealed this

Eighteenth Day of February, 2003

A handwritten signature in black ink, appearing to read "James E. Rogan", written over a horizontal line.

JAMES E. ROGAN  
*Director of the United States Patent and Trademark Office*