

2. Upon information and belief, defendant IBM Corporation is a company organized and existing under the laws of the State of New York, having its principal place of business at 1 New Orchard Road, Armonk, New York 10504-1783.

JURISDICTION AND VENUE

3. This action arises under the patent laws of the United States, 35 U.S.C. §§ 1 *et seq.* This Court has jurisdiction over the subject matter of this action pursuant to 28 U.S.C. §§ 1331 and 1338(a).

4. Venue is proper in this district pursuant to 28 U.S.C. §§ 1391(b), 1391(c), and 1400(b).

5. Personal jurisdiction over defendant comports with the United States Constitution because defendant has a regular and established place of business and is contributing and/or committing the acts of patent infringement alleged in this Complaint in this district.

CLAIM FOR RELIEF

Patent Infringement Of U.S. Patent No. 7,197,547

6. On March 27, 2007, United States Patent No. 7,197,547 ("the '547 Patent"), entitled "Load Balancing Technique Implemented in a Data Network Device Utilizing a Data Cache," was duly and lawfully issued based upon an application filed by the inventors Andrew Karl Miller, Jack Dee Menendez, and Ajit Ramachandra Mayya. (A true and correct copy of the '547 Patent is attached as Exh. A.)

7. Plaintiff is the owner by assignment of the '547 Patent, and has the right to sue and recover damages for infringement thereof.

8. Upon information and belief, defendant is committing acts of direct and indirect infringement of at least claims 1 and 18 of the '547 Patent by the manufacture, use, offer for and/or sale in the United States of systems such as the IBM Websphere Commerce system which are load balanced server farm systems for effecting electronic commerce over a data network and which practice a method for effecting electronic commerce over a data network.

9. Upon information and belief, after a reasonable opportunity for further investigation or discovery, the evidence will show the acts of infringement by defendant have been with the knowledge of the '547 Patent and are willful and deliberate, thus rendering this action against defendant "exceptional" as that is employed in 35 U.S.C. § 285.

10. The acts of infringement by defendant will continue unless enjoined by this Court.

11. Plaintiff has been and will continue to be irreparably harmed and damaged by the acts of infringement of the '547 Patent by defendant and has no adequate remedy at law.

PRAYER FOR RELIEF

WHEREFORE, plaintiff prays for the following relief:

- A. An order adjudging defendant to have infringed U.S. Patent No. 7,197,547;
- B. An order enjoining defendant together with its officers, agents, servants, employees, and attorneys, and upon those persons in active concert or participation with

them who receive actual notice of this order by personal service or otherwise, from infringing U.S. Patent No. 7,197,547;

- C. An award of compensatory damages trebled as provided by 35 U.S.C. § 284, together with interests and costs;
- D. An award of reasonable attorney fees as provided by 35 U.S.C. § 285; and
- E. Such other and further relief as this Court may deem necessary and just.

JURY DEMAND

Pursuant to Fed. R. Civ. P. 38(b), plaintiff hereby demands a trial by a jury on all issues so triable.

Respectfully submitted,

LERNER, DAVID, LITTENBERG,
KRUMHOLZ & MENTLIK, LLP

Dated: Jan. 18, 2012

By: 

Orville R. Cockings

Tel: 908.654.5000

E-mail: ockings@ldlkm.com

litigation@ldlkm.com

CERTIFICATION PURSUANT TO LOCAL CIVIL RULE 1.6(a)

The undersigned hereby certifies, pursuant to Local Civil Rule 1.6(a), that with respect to the matter in controversy herein, neither plaintiff IP Venture, Inc. nor plaintiff IP Venture, Inc.'s attorney is aware of any other action pending in any court, or of any pending arbitration or administrative proceeding, to which this matter is subject.

Dated: Jan. 18, 2012

LERNER, DAVID, LITTENBERG,
KRUMHOLZ & MENTLIK, LLP
Attorneys for Plaintiff IP Venture, Inc.

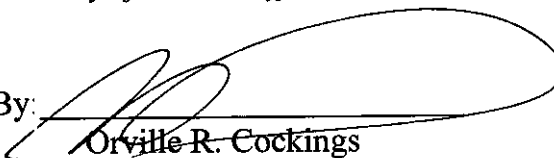
By: 
Orville R. Cockings
Tel: 908.654.5000
E-mail: ocockings@ldlkm.com
litigation@ldlkm.com

EXHIBIT A



US007197547B1

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

(10) **Patent No.:** US 7,197,547 B1
 (45) **Date of Patent:** Mar. 27, 2007

(54) **LOAD BALANCING TECHNIQUE
 IMPLEMENTED IN A DATA NETWORK
 DEVICE UTILIZING A DATA CACHE**

(76) Inventors: **Andrew Karl Miller**, 596 Panchita Way, Los Altos, CA (US) 94022; **Jack Dee Menendez**, 925 Eichler Way, Mt. View, CA (US) 94040; **Ajit Ramachandra Mayya**, 3575 Lomond Ct., Saratoga, CA (US) 94070

(*) 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.: 09/568,823

(22) Filed: May 10, 2000

Related U.S. Application Data

(60) Provisional application No. 60/133,646, filed on May 11, 1999.

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** 709/223; 718/105

(58) **Field of Classification Search** 709/223,
 709/224, 229, 105; 718/105

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

2,781,643 A	2/1957	Fairweather
3,406,532 A	10/1968	Rownd et al.
3,670,867 A	6/1972	Traube
4,213,310 A	7/1980	Buss
4,455,453 A	6/1984	Parasekvakos et al.

(Continued)

FOREIGN PATENT DOCUMENTS

FR	2696722	4/1994
GB	2 265 032 A	9/1993
WO	WO99/07121	2/1999

OTHER PUBLICATIONS

Automatic ID News, "20/20 Results Achieved with Technology Trio", Sep. 1995, p. 19.

Henry Towie, "On the Fast Track with Totaltracks: UPS Deploys Mobile Date Service," Abstract No., XP-00060076, Document Delivery World, vol. 9, No. 3, 1993, pp. 30-31.


(Continued)

Primary Examiner—Ario Etienne
Assistant Examiner—Barbara Burgess

(57) **ABSTRACT**

A technique for implementing a load balanced server farm system is described which may be used for effecting electronic commerce over a data network. The system comprises a load balancing system and a plurality of servers in communication with the load balancing system. Each of the plurality of servers may include a respective data cache for storing state information relating to client session transactions conducted between the server and a particular client. The load balancing system is configured to select, using a load balancing protocol, an available first server from the plurality of servers to process an initial packet received from a source device such as, for example, a client machine of a customer. The load balancing system is also configured to route subsequent packets received from the source device to the first server. In this way, a "stickiness" scheme may be implemented in the server farm system whereby, once an electronic commerce session has been initiated between the first server and the source device, the first server may handle all subsequent requests from the source device in order to make optimal use of the state data stored in the first server's data cache. Before generating its response, the first server may verify that the state information relating to a specific client session stored in the data cache is up-to-date. If the first server determines that the state information stored in the data cache is not up-to-date, then the first server may be configured to retrieve the desired up-to-date state information from a database which is configured to store all state information relating to client sessions which have been initiated with the server farm system.

35 Claims, 9 Drawing Sheets



US 7,197,547 B1

Page 2

U.S. PATENT DOCUMENTS

4,656,591 A	4/1987	Goldberg	6,292,784 B1	9/2001	Martin et al.
4,799,156 A	1/1989	Shavit et al.	6,324,520 B1	11/2001	Walker et al.
4,887,208 A	12/1989	Schneider et al.	6,332,334 B1	12/2001	Faryabi
4,936,738 A	6/1990	Brennan	6,341,269 B1	1/2002	Dulaney et al.
5,038,283 A	8/1991	Caveney	6,343,275 B1	1/2002	Wong
5,093,794 A	3/1992	Howie et al.	6,397,246 B1	5/2002	Wolfe
5,105,627 A	4/1992	Kurita	6,405,173 B1	6/2002	Honarvar et al.
5,122,959 A	6/1992	Nathanson et al.	6,424,947 B1	7/2002	Tsuria et al.
5,235,819 A	8/1993	Bruce	6,445,976 B1	9/2002	Ostro
5,237,158 A	8/1993	Kern et al.	6,453,306 B1	9/2002	Quelene
5,246,332 A	9/1993	Bernard	6,463,345 B1	10/2002	Peachey-Kountz et al.
5,265,006 A	11/1993	Asthana	6,463,420 B1	10/2002	Guidice et al.
5,272,638 A	12/1993	Martin et al.	6,450,567 B1	12/2002	Gregory
5,273,392 A	12/1993	Bernard	6,456,205 B1	12/2002	White et al.
5,322,406 A	6/1994	Pippin et al.	6,505,093 B1	1/2003	Thatcher et al.
5,363,310 A	11/1994	Haj-Ali-Ahmadi et al.	6,505,171 B1	1/2003	Cohen et al.
5,395,206 A	3/1995	Cerny, Jr.	6,526,392 B1	2/2003	Dietrich et al.
5,428,546 A	6/1995	Shah et al.	6,520,518 B1	3/2003	Krichilsky et al.
5,533,361 A	7/1996	Halpern	6,567,786 B1	5/2003	Bibelnieks et al.
5,548,518 A	8/1996	Dietrich et al.	6,511,213 B1	5/2003	Altendahl et al.
5,593,269 A	1/1997	Bernard	6,518,005 B1	6/2003	Lesaint et al.
5,615,121 A	3/1997	Babayev et al.	6,598,027 B1	7/2003	Breen, Jr. et al.
5,666,493 A	9/1997	Wojcik et al.	6,512,127 B1	9/2003	Klots et al.
5,694,551 A	12/1997	Doyle et al.	6,514,726 B1	11/2003	Hanzek
5,712,989 A	1/1998	Johnson et al.	6,697,964 B1	2/2004	Doddrill et al.
5,758,313 A	5/1998	Shah et al.	6,711,995 B1	5/2004	Chen et al.
5,758,328 A	5/1998	Giovannoli	6,748,418 B1	6/2004	Yoshida et al.
5,761,673 A	6/1998	Bookman et al.	6,763,496 B1	7/2004	Hennings et al.
5,768,139 A	6/1998	Pippin et al.	6,862,572 B1	3/2005	de Sylva
H1743 H	8/1998	Graves et al.	6,970,837 B1	11/2005	Walker et al.
5,809,479 A	9/1998	Martin et al.	6,990,460 B2	1/2006	Parkinson
5,826,242 A	10/1998	Montulli	2001/0037229 A1	11/2001	Jacobs et al.
5,826,825 A	10/1998	Gabriel	2001/0042021 A1	11/2001	Matsuo et al.
5,831,860 A	11/1998	Foladare et al.	2001/0047285 A1	11/2001	Borders et al.
5,832,457 A	11/1998	Cherney	2001/0047310 A1	11/2001	Russell
5,834,753 A	11/1998	Danielson et al.	2001/0049619 A1	12/2001	Powell et al.
5,835,914 A	11/1998	Brim	2001/0049672 A1	12/2001	Moore
5,839,117 A	11/1998	Cameron et al.	2002/0004766 A1	1/2002	White
5,848,395 A	12/1998	Edgar et al.	2002/0007299 A1	1/2002	Florence
5,878,401 A	3/1999	Joseph	2002/0013950 A1	1/2002	Tomsen
5,880,443 A	3/1999	McDonald et al.	2002/0038224 A1	3/2002	Bhadra
5,893,076 A	4/1999	Hafner et al.	2002/0049853 A1	4/2002	Chu et al.
5,894,554 A	4/1999	Lowery et al.	2002/0065700 A1	5/2002	Powell et al.
5,897,622 A	4/1999	Blinn et al.	2002/0088530 A1	12/2002	Wojcik et al.
5,897,629 A	4/1999	Shinagawa et al.	2002/0094087 A1	12/2002	Spiegel et al.
5,899,088 A	5/1999	Purdum	2003/0045340 A1	3/2003	Roberts
5,910,896 A	6/1999	Hahn-Carlson	2003/0079227 A1	4/2003	Knowles et al.
5,918,213 A	6/1999	Bernard et al.	2003/0133190 A1	12/2003	Jones
5,943,652 A	8/1999	Sisley et al.	2004/0136635 A1	11/2004	Publicover
5,943,841 A	8/1999	Wunscher	2005/0127580 A1	2/2005	Crici et al.
5,956,709 A	9/1999	Xue	2005/0144641 A1	6/2005	Lewis
5,963,919 A	10/1999	Brinkley et al.			
5,979,757 A	11/1999	Tracy et al.			
6,023,683 A	2/2000	Johnson et al.			
6,061,607 A	5/2000	Bradley et al.			
6,070,147 A	5/2000	Harms et al.			
6,073,108 A	6/2000	Peterson			
6,081,789 A	6/2000	Purcell			
6,083,279 A	7/2000	Cuomo et al.			
6,085,170 A	7/2000	Tsukuda			
6,101,481 A	8/2000	Miller			
6,140,922 A	10/2000	Kakou			
6,178,510 B1	1/2001	O'Connor et al.			
6,185,625 B1	2/2001	Tso et al.			
6,215,952 B1	4/2001	Yoshio et al.			
6,233,543 B1 *	5/2001	Butts et al. 703/27			
6,249,801 B1 *	6/2001	Zisapel et al. 709/105			
6,260,024 B1	7/2001	Shkedy			
6,275,812 B1	8/2001	Haq et al.			
6,289,260 B1	9/2001	Bradley et al.			

OTHER PUBLICATIONS

Hiroo Kawata, "Information Technology of Commercial Vehicles in the Japanese Parcel Service Business," Abstract No., XI-000560489, 1992, pp. 371-382.

Koster, Rene de, "Routing Orderpickers in a Warehouse: A Comparison Between Optimal and Heuristic Solutions," IIE Transactions, vol. 30, NO. 5, p 469, May 1998.

Maloney, David, "The New Corner Drugstore", May 1, 2000, Modern Materials Handling, vol. 55, No. 5, p. 58.

PC Foods, "Customer Service Agreement," printed form website: <http://www.pcfoods.com>, Abstract No., XP-002245026, 1999, pp. 1-2.

Takashi Sekita, "The Physical Distribution Information Network in the Home-Delivery Business," Japan Computer Quarterly, Abstract No., XP-00.431194, 1990, pp. 23-32.

US 7,197,547 B1

Page 3

The Impact of Electronic Data Interchange on Competitiveness in Retail Supply Chain, Brian Fynes et al., IBAR v14n2 pp. 16-2 1993.

Van Den Berg, Jeroen, P, "A Literature Survey on Planning and Control of Warehousing Systems", IIE Transactions vol. 31, No. 3, p. 751, Aug. 1999.

Vass et al., "The World Wide Web—Everything you (n)ever wanted to know about its server", IEEE, Oct./Nov. 1998, pp. 33-37.

Wilson, Joe, "Selecting Warehouse Management Software (WMS) for Food Distribution Operations", Frozen Food Digest, Oct. 1998, vol. 14, NO. 1, p. 18.

Wunnava et al., "Interactive Multimedia on the World Wide Web", IEEE, Aug. 1999, pp. 110-115.

U.S. Appl. No. 09/568,570, filed May 10, 2000.

U.S. Appl. No. 09/568,571, filed May 10, 2000.

U.S. Appl. No. 09/568,572, filed May 10, 2000.

U.S. Appl. No. 09/568,603, filed May 10, 2000.

U.S. Appl. No. 09/568,613, filed May 10, 2000.

U.S. Appl. No. 09/568,614, filed May 10, 2000.

U.S. Appl. No. 09/620,199, filed Jul. 20, 2000.

U.S. Appl. No. 09/750,385, filed Dec. 27, 2000.

U.S. Appl. No. 09/792,400, filed Feb. 22, 2001.

U.S. Appl. No. 09/813,235, filed Mar. 19, 2001.

Hyten, Todd, "Stop & Shop mulls online grocery store", Boston Business Journal (Boston, MA, US), vol.: 16, No. 6, p. 1, Mar. 22, 1996.

Pearce, Michael R. "From carts to clicks", Ivey Business Quarterly, vol. 63, No. 1, p. 69-71, Autumn 1998.

"Peapod Interactive Grocery Shopping and Delivery Service Now Delivers Via the Internet", Press Release, peapod.com, Apr. 22, 1996, pp. 1-2.

Worth Vren, Jr., Fort Worth Star-Telegram Texas, "Albertson's Expects Online Grocery Shopping To Boom", KRFBN Knight-Ridder Tribune Business News (Fort Worth Star-Telegram Texas), Nov. 9, 1998.

www.peapod.com, including introduction to Peapod: How Peapod Works; Peapod: Choosing a Delivery Time; Peapod: Sending Your Order; Retrieved from Internet Archive (web.archive.org) on Jul. 23, 2006, alleged date Nov. 13, 1996, pp. 1-9.

Norton, Tim R., "End-To-End Response-Time: Where to Measure?", Computer Measurement Group Conference Proceedings, CMG99, Session 423, Dec. 1999, pp. 1-9.

Smith et al., "Management of Multi-Item Retail Inventory Systems with Demand Substitution", Operations Research, vol. 48, No. 1, Jan.-Feb., pp. 50-64.

Anupindi et al., "Estimation of Consumer Demand with Stock-Out Based Substitution: An Application to Vending Machine Product", Marketing Science, vol. 17, No. 4, 1998, pp. 406-423.

* cited by examiner




Fig. 1




Fig. 2




Fig. 3

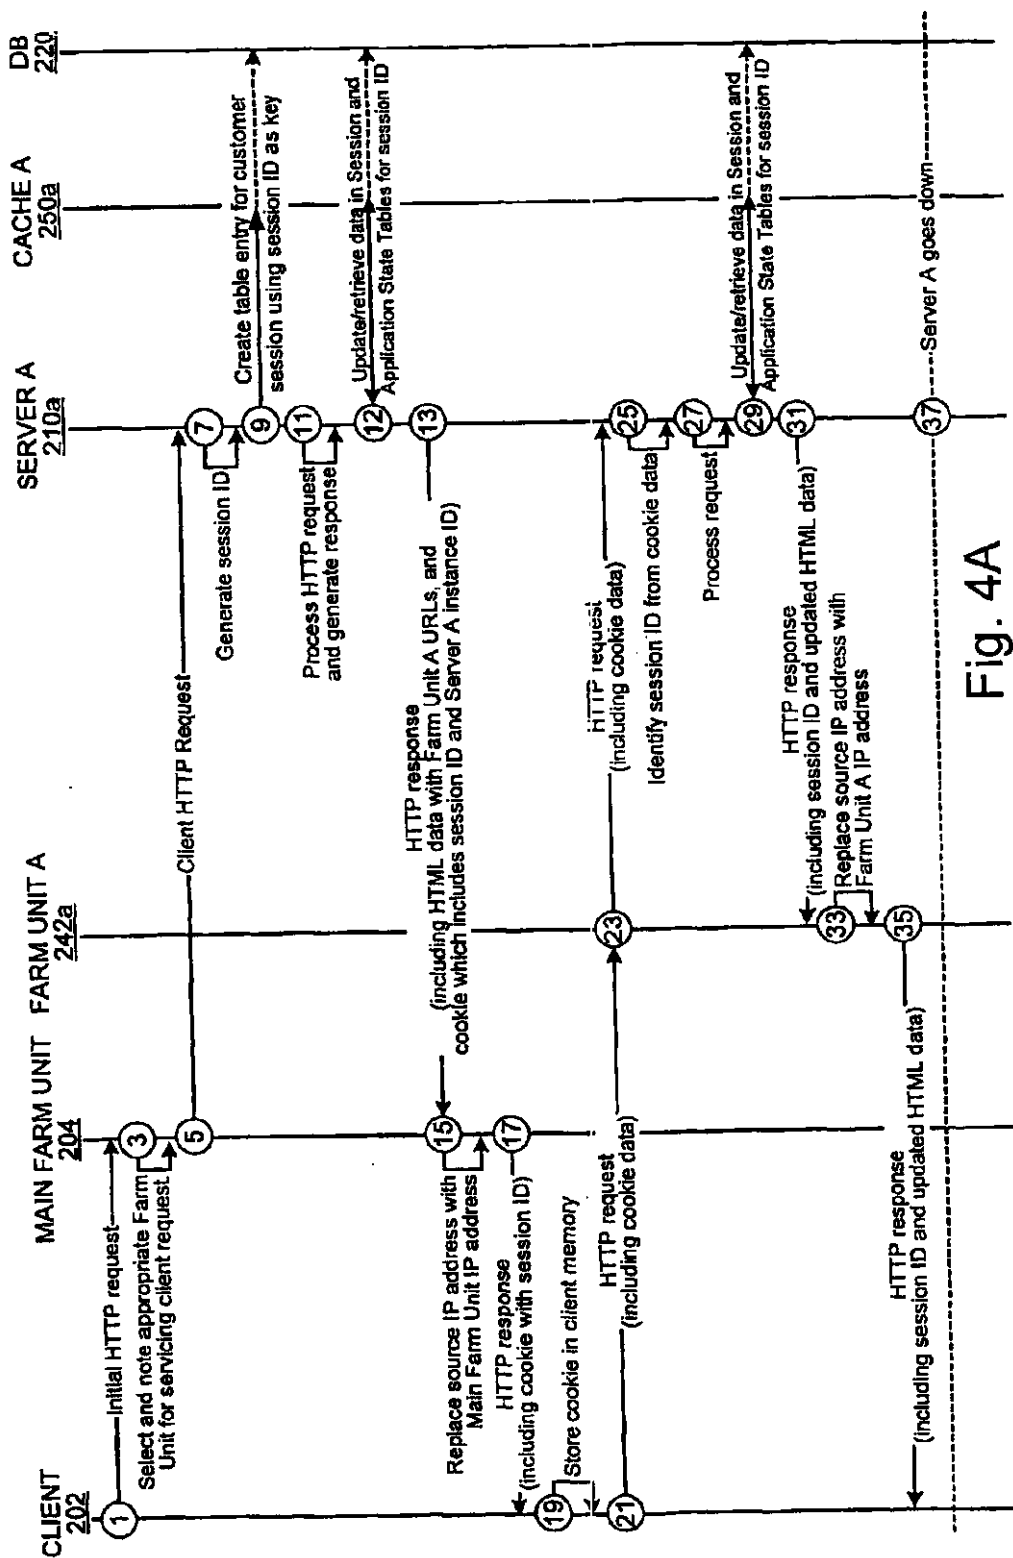


Fig. 4A

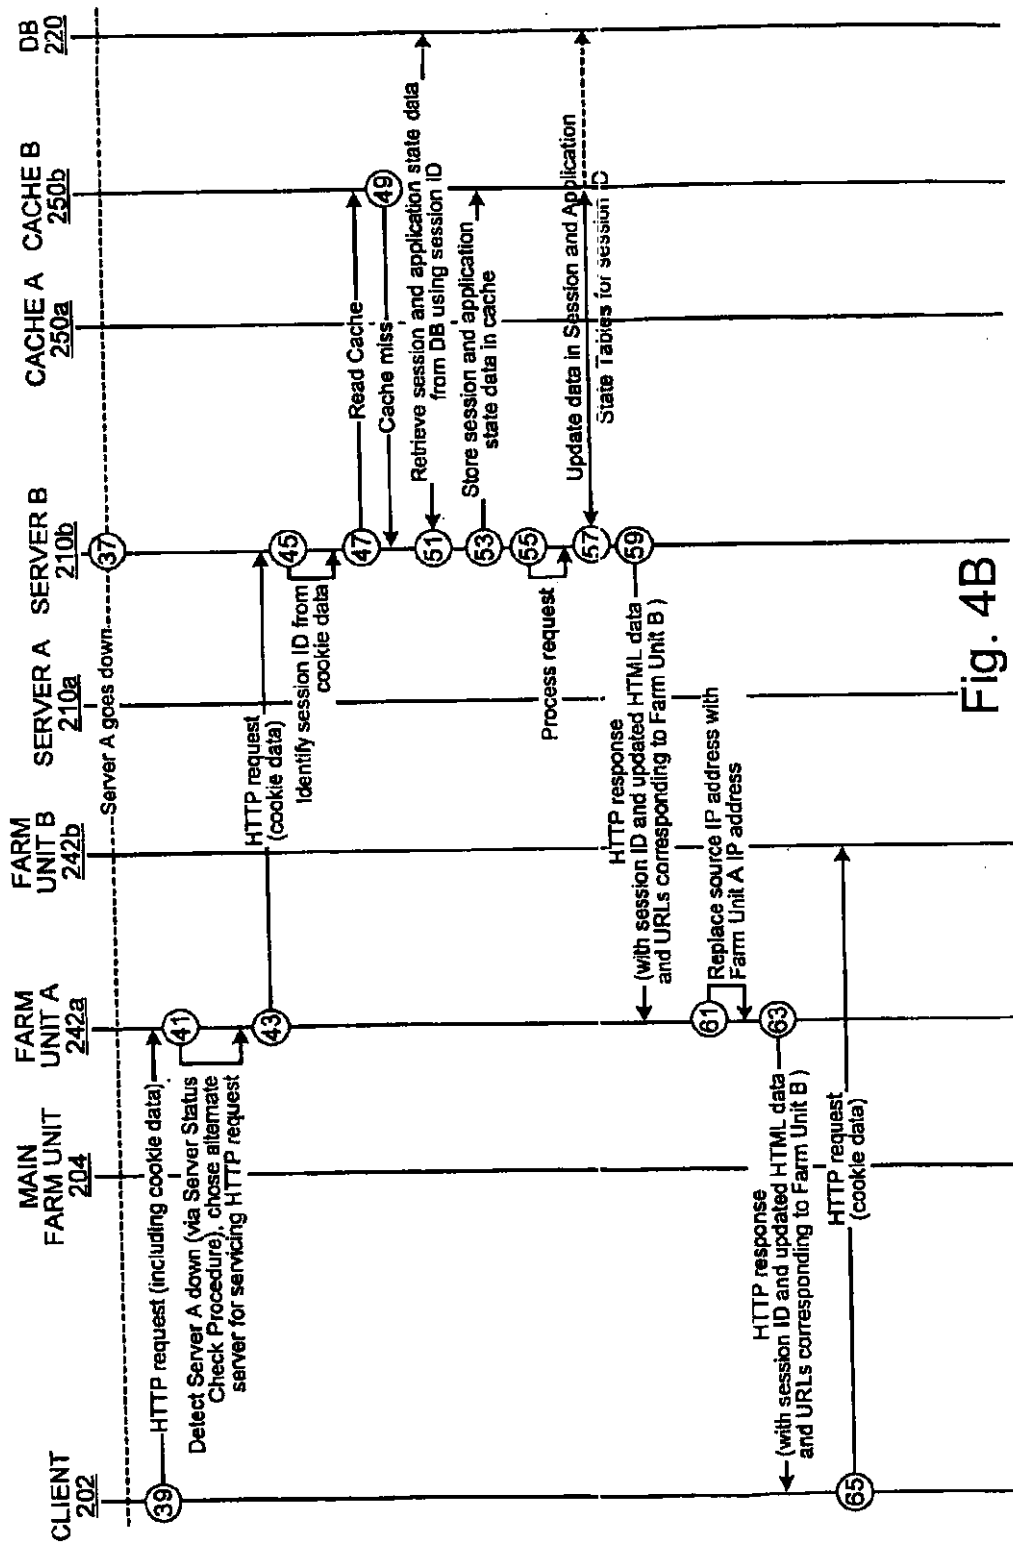


Fig. 4B

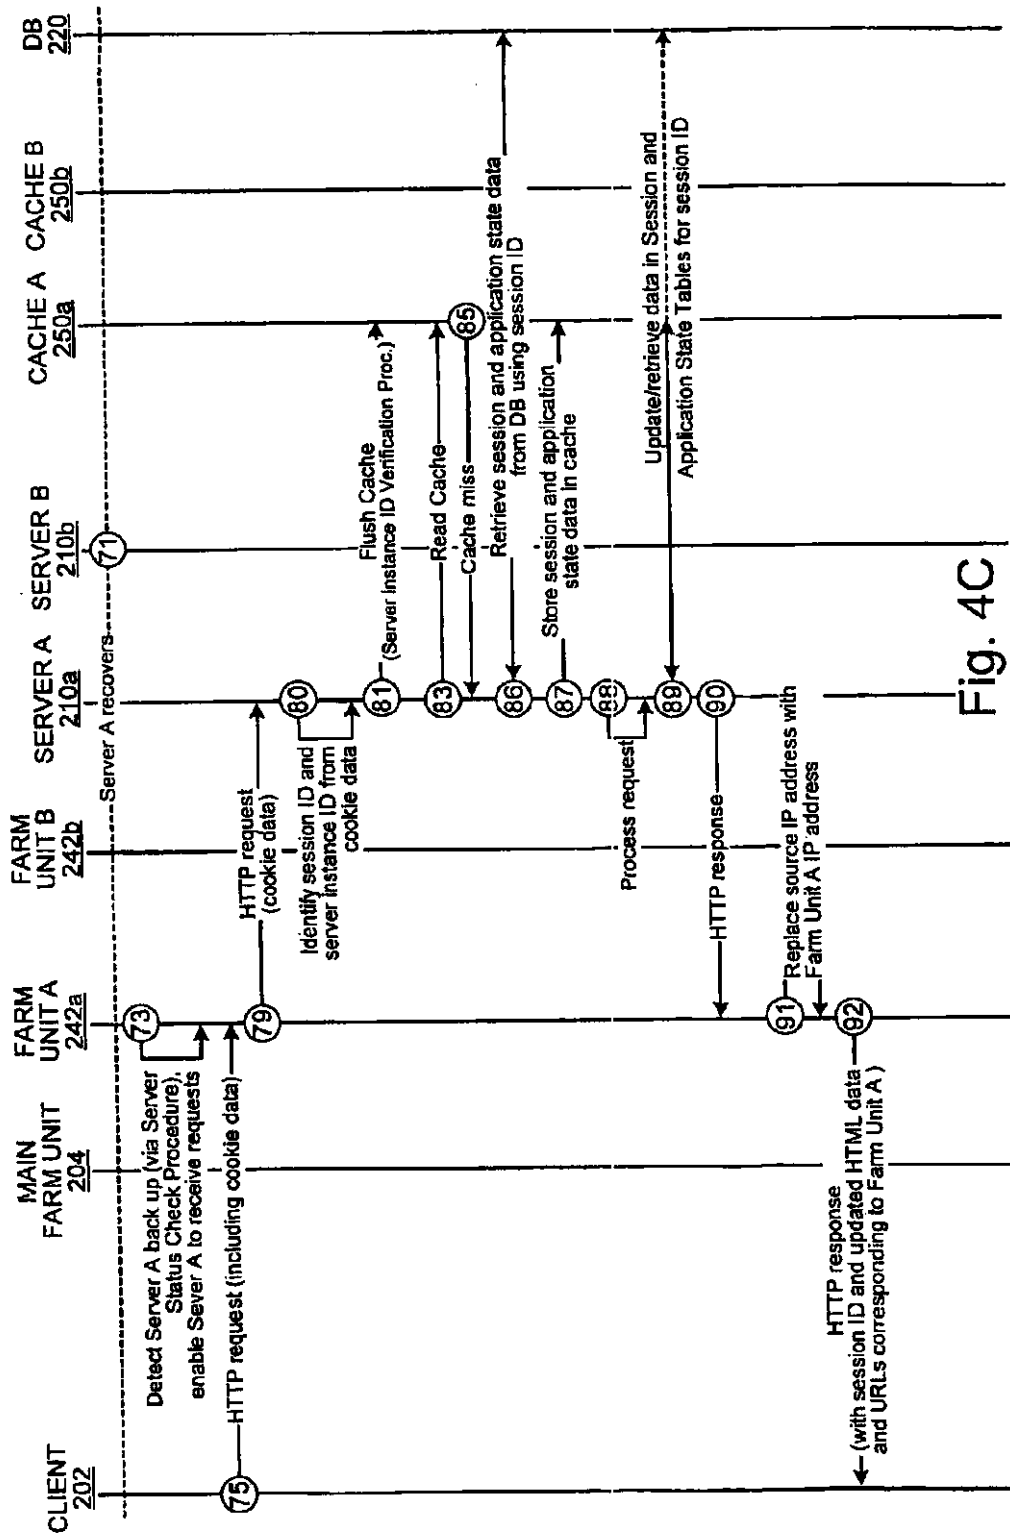


Fig. 4C

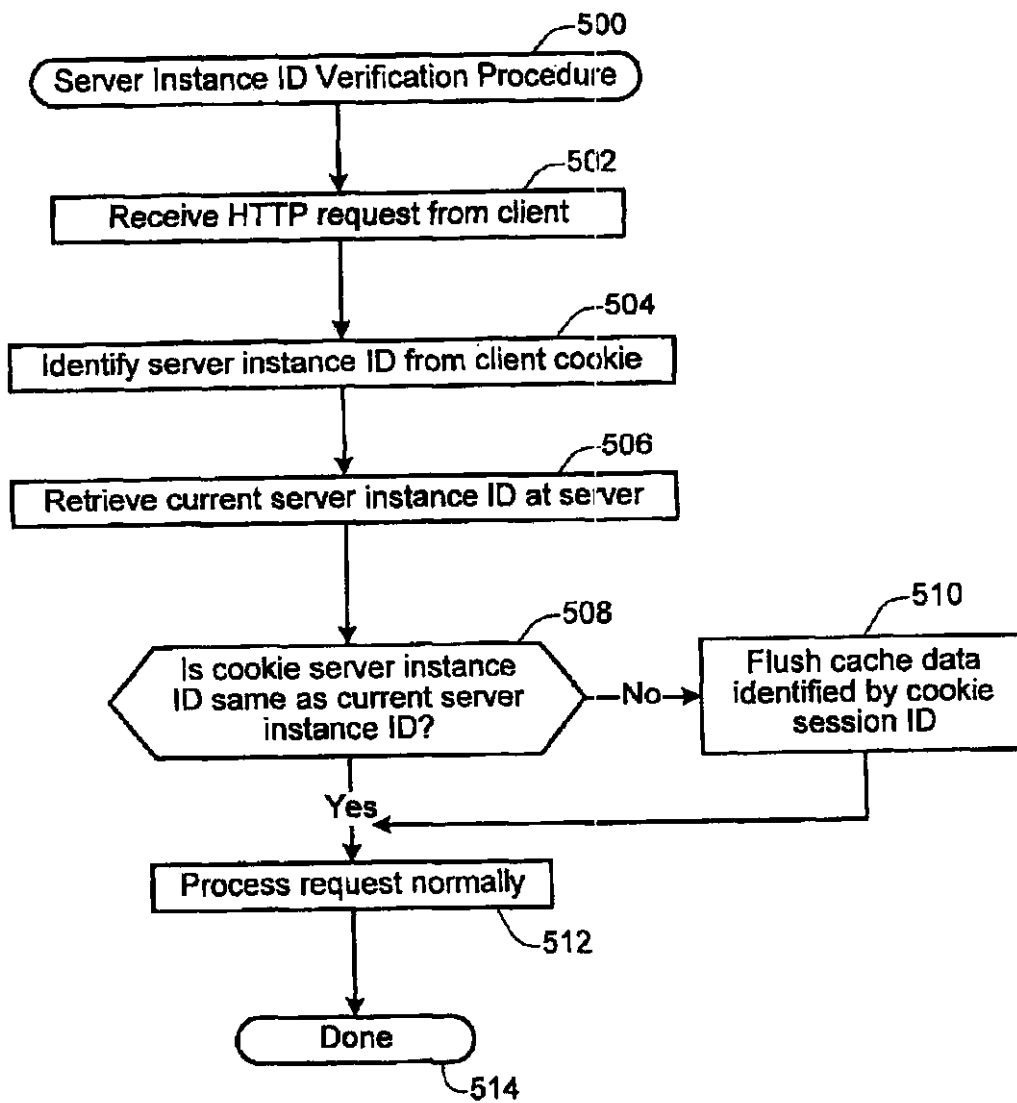


Fig. 5

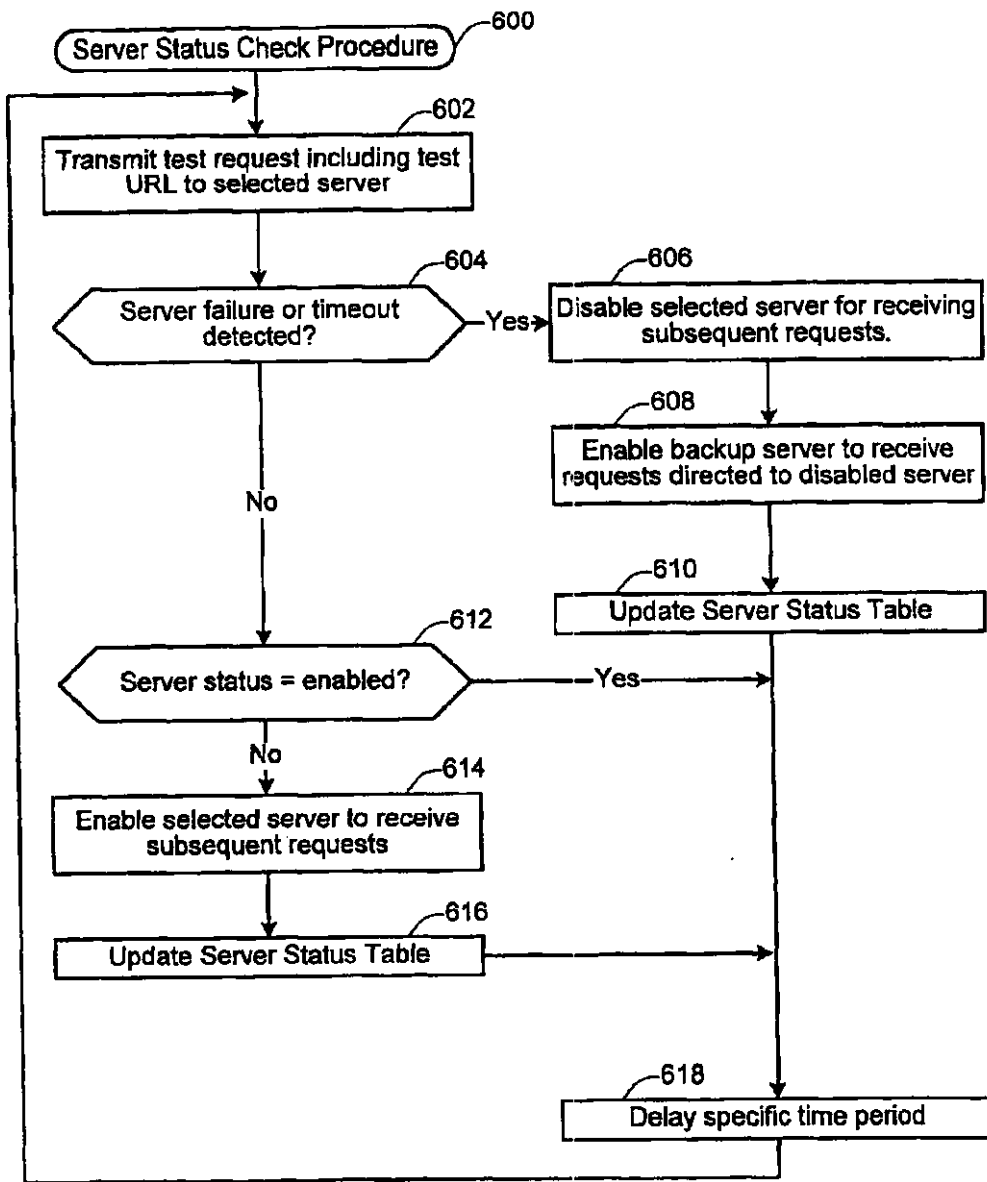


Fig. 6




Figure 7

US 7,197,547 B1

1

LOAD BALANCING TECHNIQUE IMPLEMENTED IN A DATA NETWORK DEVICE UTILIZING A DATA CACHE

RELATED APPLICATION DATA

The present application claims priority under 35 USC 119(e) from U.S. Provisional Pat. application Ser. No. 60/133,646 entitled ELECTRONIC COMMERCE ENABLED DELIVERY SYSTEM AND METHOD filed May 11, 1999, the entirety of which is incorporated herein by reference for all purposes.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to data networks, and more specifically to a load balancing technique implemented in a data network device utilizing a data cache

2. Description of the Related Art

Content providers on the World Wide Web are willing to pay a great deal of money to ensure that their information is provided quickly and efficiently to any given client or user. Recently, there has been a great deal of research effort directed at reducing network congestion and increasing server response time. One particular area which has received much attention relates to load balancing techniques for network traffic.

FIG. 1 shows a schematic block diagram of a conventional load balancing implementation which may be used to increase server response time to a given user. In the example of FIG. 1, a client or user 102 desires to access a specific web page associated with a specific URL such as, for example, www.yahoo.com. As shown in FIG. 1, the content provider associated with the desired URL has implemented a load balanced host server system 130 which includes a load balancing device 104, a farm of web servers 110, and a centralized database 120. Each server 110a, 110b, 110n of the server farm 110 includes identical content corresponding to the content provider's entire web site. When the user 102 attempts to access the content provider's web site, via gateway device 106 and the Internet 108, the user is first directed to the load balancing device 104. The load balancing device then dynamically determines which server of the server farm 110 is the least busy, and forwards the client request to the identified server. Thus, for example, if the load balancing device 104 determines that Server A 110a is the least busy, the load balancing device will forward the client request (e.g., client HTTP request) to Server 110a. Server 110a then generates a response to the client request, and transmits the response to load balancing device 104, where it is then forwarded to the client 102. Each time the client transmits an HTTP request to the load balancing device 104, the load balancing device selects an appropriate (e.g., least busy) server of the server farm 110 to respond to the client request. Moreover, since the content on each server is identical, more than one server may be used for responding to requests from a particular client. Further, it will be appreciated that the host server system 130 also provides fail over protection by way of multiple redundant servers (e.g., 110a, 110b, 110n).

Implementing State In TCP

As commonly known to one having ordinary skill in the art, TCP (Transmission Control Protocol) is a stateless protocol. Thus, in order to implement state over TCP/IP, content providers conventionally use a "cookie file" file to maintain state information for a given client. Typically, the cookie file will be transmitted to the client where it is then

2

stored on the client machine. When the client machine accesses a particular web server, the web server retrieves the appropriate cookie file data from the client machine in order to properly respond to the client. According to one conventional technique, the client cookie file will include all necessary state information relating to the client's current session with the web server. This implementation may be useful in situations where there is a relatively small amount of state information to be stored in the cookie file. However, problems may be encountered when there is a relatively large amount of state or other information to be stored in the cookie file. For example, if the user accesses an electronic commerce site such as, for example, an on-line grocery store, the user may select dozens or even hundreds of items to add to his or her electronic shopping cart. Each time the user adds a new item to the electronic shopping cart, the state information for that user needs to be updated to include the new item. Moreover, a new cookie file which includes the updated state information (including the updated contents of the user's shopping cart) must be transmitted back to the user's computer. This may result in a significant decrease in response time as experienced by the user. Further, as the user continues to add new items to the shopping cart, the relative response time experienced by the user may continue to decrease.

A second conventional technique for implementing state over TCP/IP provides that the state information relating to a particular client session be stored in a database on the host server system. According to this later technique, when the client first accesses the host server system 130, a session ID is generated for that client session. The session ID is then stored in a cookie file on the client machine 102. The state information corresponding to that client session is stored on the host server database 120, and may be accessed using the session ID. Thereafter, during the client session, each time the client accesses the host server system 130, the assigned host server from the server farm 110 will retrieve the session ID data from the cookie file stored on the client machine, and, using the session ID, will retrieve the appropriate state information from database 120. Thus, according to this technique, when a client adds a new item to his or her shopping cart, for example, the assigned host server will update the client's state information stored on database 120.

Although the later-described technique for implementing state over TCP/IP reduces the amount of data to be written to the client cookie file, it necessarily involves accessing the database 120 each time one of the farm servers 110 desires to read or write state information relating to a particular client session ID. This results in a decreased response time from the host server system 130, as experienced by the user 102. Accordingly, there exists a continual need to improve upon network load balancing and fail over protection techniques.

SUMMARY OF THE INVENTION

According to specific embodiments of the present invention, a load balanced server farm system is provided for effecting electronic commerce over a data network. The system comprises a load balancing system and a plurality of servers in communication with the load balancing system. Each of the plurality of servers may include a respective data cache for storing state information relating to client session transactions conducted between the server and a particular client. The load balancing system is configured to select, using a load balancing protocol, an available first server from the plurality of servers to process an initial packet received from a source device such as, for example, a client

US 7,197,547 B1

3

machine of a customer. The load balancing system is also configured to route subsequent packets received from the source device to the first server. In this way, a "stickiness" scheme may be implemented in the server farm system whereby, once an electronic commerce session has been initiated between the first server and the source device, the first server may handle all subsequent requests from the source device in order to make optimal use of the state data stored in the first server's data cache. One technique implementing the above-described "stickiness" scheme is to configure the content on each of the plurality of servers to include a respective plurality of Uniform Resource Locators (URLs) which correspond to addresses for accessing information specific to the server on which the URL resides.

An additional aspect of the present invention provides that one or more of the subsequent packets received from the source device may include a session ID corresponding to an electronic commerce session initiated at the server farm system for the source device. The first server is configured to access, from the data cache, state information relating to the electronic commerce session associated with the specified session ID. The first server is also configured to generate a response to a subsequent request packet received from the source device using state information retrieved from the data cache. Before generating its response to the subsequent request packet, the first server may verify that the state information relating to the client session stored in the data cache is up-to-date. If the first server determines that the state information stored in the data cache is not up-to-date, then the first server may be configured to retrieve the desired up-to-date state information from a database which is configured to store all state information relating to client sessions which have been initiated with the server farm system.

An alternate embodiment of the present invention is directed to a system for effecting electronic commerce over a data network. The system comprises means for receiving an initial request packet from a source device, and means for performing a load balancing protocol, wherein the initial packet from the source device is assigned to a first server of a load balanced server farm. The load balanced server farm system may include a plurality of different servers. The first server may comprise a first data cache. The system further comprises means, at the first server, for generating a first response to the initial request packet; means for transmitting the first response to the source device; and means for causing subsequent packets received from the source device to be routed to the first server.

Other embodiments of the present invention are directed to a method or computer program product for effecting electronic commerce over a data network. An initial request packet from a source device is received. A load balancing procedure is then performed, wherein the initial packet from the source device is assigned to a first server of a load balanced server farm system. The load balance server farm system includes a plurality of different servers, each of which includes a respective data cache. A first response to the initial request packet is generated at the first server. The first response is then transmitted to the source device. The response transmitted to the source device causes subsequent packets received from the source device to be routed to the first server.

An alternate embodiment of the present invention is directed to a system for implementing fail over protection of a load balanced server farm system connected to a data network. The system comprises a load balancing system including a main server farm unit and a plurality of subordinate server farm units. The system further includes a

4

plurality of servers in communication with the load balancing system, wherein each server of the plurality of servers is associated with a respective subordinate server farm unit. The plurality of subordinate server farm units comprises a first server farm unit which includes a first server. The plurality of subordinate server farm units also comprises a second server farm unit which includes a second server. The system is configured to cause a first plurality of packets received from a source device to be routed to the first server farm unit while a failure at the first server is not detected. The system is further configured to cause a second plurality of packets received from the source device to be routed to the second server farm unit while a failure at the first server is detected. An additional aspect of this embodiment provides that each of the servers is configured to generate respective responses to client requests, wherein at least a portion of the responses includes URLs for accessing additional information from the specific server which generated the response.

Further embodiments of the present invention are directed to a method and computer program product for implementing fail over protection of a load balanced server farm system connected to a data network. A first request packet is received from a source device, the first request packet includes session ID information for identifying an initiated communication session between the source device and a first server of the server farm system. A failure is then detected at the first server. The first request packet is then routed to a second server selected from the server farm system in response to detecting the first server failure. A first response to the first request packet is then generated. The first response includes at least one URL for accessing information via the second server. A response packet which includes the first response is then transmitted to the source device. The response packet includes a source IP address corresponding to the first server. Additional objects, features and advantages of the various aspects of the present invention will become apparent from the following description of its preferred embodiments, which description should be taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a schematic block diagram of a conventional load balancing implementation.

FIG. 2 shows a block diagram of a specific embodiment of a data network 200 which may be used for implementing the technique of the present invention.

FIG. 3 shows a block diagram of a specific implementation for storing customer session and application state data in data cache 250a

FIGS. 4A, 4B and 4C illustrate data flow diagrams corresponding to a specific implementation of the present invention.

FIG. 5 shows a flow diagram of a Server Instance ID Verification Procedure 500 in accordance with a specific embodiment of the present invention.

FIG. 6 shows a flow diagram of a Server Status Check Procedure 600 in accordance with a specific embodiment of the present invention.

FIG. 7 shows a specific embodiment of a server device 700 suitable for implementing a server of present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIG. 2 shows a block diagram of a specific embodiment of a data network 200 which may be used for implementing

US 7,197,547 B1

5

the technique of the present invention. According to the embodiment of FIG. 2, the host server system 230 may be implemented as a load-balanced server farm which includes a load balancing device 240 and a plurality of servers 210. The host server system 230 of the present invention may be used, for example, to implement the Webstore Subsystem described in U.S. Pat. application Ser. No. 09/568,603 for INTEGRATED SYSTEM FOR ORDERING, FULFILLMENT, AND DELIVERY OF CONSUMER PRODUCTS USING A DATA NETWORK, filed simultaneously herewith, which is incorporated herein by reference in its entirety for all purposes.

According to a specific implementation, the load balancing device 240 may be implemented using a conventional load balancing device such as, for example, the RND device manufactured by Radware, Inc., of Mahwah, N.J.

As shown in FIG. 2, the load balancing device 240 may be configured to include a Main Farm Unit 204 and a plurality of subordinate Farm Units 242. According to a specific embodiment, the Main Farm Unit 204 and plurality of subordinate Farm Units 242 may be implemented as logical devices using hardware and/or software. For example, in one embodiment, the Main Farm Unit 204 may be implemented as a logical device on the load balancing device 242. Similarly, one or more of the subordinate Farm Units (e.g., 242a, 242b, 242n) may also be implemented as logical devices on the load balancing device 240. Alternatively, it will be appreciated that the Main Farm Unit 204 and plurality of subordinate Farm Units 242 may each be implemented on separate physical devices which are part of the same computer system or network.

As shown in the embodiment of FIG. 2, each of the plurality of subordinate Farm Units 242 has associated with it one or more respective host servers, collectively identified by reference number 210 of FIG. 2. Thus, for example, as shown in FIG. 2, Farm Unit A 242a is associated with a host Server A 210a, Farm Unit B 242b is associated with a host Server B 210b, and Farm Unit N 242n is associated with a host Server N 210n. It will be appreciated that, in alternative embodiments, one or more of the subordinate Farm Units 242 may be associated a respective plurality of farm servers. For example, Farm Unit A 242a may be configured to be associated with a plurality of host servers assigned to that particular subordinate Farm Unit.

Additionally, as shown in FIG. 2, each of the plurality of host servers 210 is configured to be associated with a respective data cache 250. Thus, for example, Server A 210a is associated with Data Cache A 250a, Server B 210b is associated with data cache B 250b, and Server N 210n is associated with data cache N 250n. According to an alternate embodiment where one or more subordinate Farm Units each is associated with a plurality of servers, each of the plurality of servers associated with a particular subordinate Farm Unit may be configured to access a common data cache for caching data relating to client requests which are processed by any of the servers belonging to the subordinate Farm Unit.

According to at least one embodiment of the present invention, each of the plurality of data caches 250 may be used to store state information for client session transactions which are processed by the host server associated with that data cache. Thus, for example, state data for client sessions which occur at Server A 210a may be stored in cache A 250a, and state data relating to client sessions which occur on Server B 210b may be stored in cache B 250b, etc. According to a specific embodiment, the state data stored in

6

a data cache may include session state data (e.g., user ID, login data, etc.) and application state data (e.g., electronic shopping cart data). As shown in FIG. 2, the server system 230 may also include a database 220 which may be used for storing session and/or application state data similar to that stored on each of the plurality of data caches 250. According to a specific implementation, the database 220 may be implemented using persistent memory, whereas one or more data caches may be implemented using volatile memory.

The technique of the present invention takes advantage of the fact that the access time for accessing data in a data cache is significantly less than that associated with accessing data from persistent memory devices. Accordingly, one advantage of the technique of the present invention is that the response time for accessing the server system of the present invention is significantly faster than that of conventional server systems such as the system shown in FIG. 1 of the drawings.

Moreover, as explained in greater detail below, the technique of the present invention solves a number of additional problems to be overcome in order for a host server of a load-balanced server farm to properly maintain client session state information on a local data cache if that server is to be used for serving HTTP requests from any given client on the World Wide Web.

In order to gain a better understanding of the problems involved with using a local data cache to maintain state information for a web farm server, it is helpful to review current techniques used by service providers to allow their clients to access web servers via the World Wide Web. Typically, most internet service providers (ISPs) provide specific gateway routers for providing client access to the World Wide Web. This is shown, for example, in FIG. 1 of the drawings. When the client 102 wishes to access a specific web site (associated with a specific host server system 130), the ISP client first sends an HTTP request to a specific gateway router 106, which then forwards the request, via Internet 108, to the specified host server system. The HTTP request packet transmitted to the host server system will include a source IP address in the header portion which identifies the source device which sent the packet, such as, for example the client machine 102 or a dedicated proxy server which, for example, may reside at gateway device 106. However, some ISPs such as, for example, America On-Line, Inc. (herein referred to as AOL) use a load-balanced farm of proxy servers to enable their clients to access the World Wide Web. In this situation, a plurality of different proxy servers may transmit packets which originate from a specific AOL client. Accordingly, it is typically the case that packets which are received at a host server system from a specific AOL client will include different source IP addresses, depending upon the particular proxy server which sent the packet.

In applying this knowledge to an example using the system of FIG. 1, it is assumed that the client machine 102 has initiated a session with host server 110a. Further, it is assumed that the host server 110a is associated with a data cache for storing state information relating to the session with client 102. When the load-balancing device 104 receives a subsequent HTTP request packet from client 102, according to the conventional load-balancing technique (described previously with respect to FIG. 1) the load balancing device may farm the request to a different host server such as, for example, Server B 110b. In this situation, Server B would be unaware of the state data cached in the memory of Server A. This will most likely result in Server B responding inappropriately to the client request.

US 7,197,547 B1

7

In order to remedy this situation, an additional step should preferably be performed by the load balancing device 104, wherein the device maintains a list of current sessions initiated with each server in the server farm 110. When the HTTP request packet from client 102 is then received at the load balancing device, the load balancing device may then inspect the source IP address of the received packet and use this address to identify the particular host server (e.g., Server A 110a) for which a session with the identified client has already been initiated. Thereafter, each time the client 102 sends a request to the host server system 130, the request will automatically be forwarded to Server A in order to utilize that session's state information which has been cached on Server A.

The problem, however, becomes more complicated when an AOL client initiates a session with the host server system 130. To illustrate this point, reference is again made to the system of FIG. 1, which now includes the improvements described in the preceding example. Additionally, it is assumed that client 102 corresponds to an AOL client. Since AOL uses a farm of proxy servers (not shown) to allow its clients to access the Internet, each request packet which is received at the load balancing device 104 from the AOL client 102 may include a different source IP address in the packet header. Accordingly, the load balancing device 104 will be unable to determine the specific host server which has already initiated a session with the AOL client. Due to the fact that there is no simple solution to this problem, most conventional load balancing and redundancy techniques resort to storing state information for a client session on a centralized database 120 which may be accessed by each of the servers in the farm.

However, contrary to conventional practices, the technique of the present invention offers a practical solution for enabling client session state information to be accessed from a data cache on a server in a load-balanced and/or redundant server farm.

FIGS. 4A, 4B and 4C illustrate data flow diagrams corresponding to a specific implementation of the present invention. The data flow diagrams of FIGS. 4A, 4B and 4C will now be described with reference to FIG. 2 of the drawings.

At (1), the client device 202 transmits an initial request to the host server system 230. In the example of FIGS. 4A, 4B, and 4C it is assumed that the client machine 202 includes a web browser which transmits HTTP requests to the host server system 230. As shown in FIG. 4A, the client request is received at Main Farm Unit 204. At (3) the Main Farm Unit 204 selects an appropriate subordinate Farm Unit (from the plurality of subordinate Farm Units 242) for servicing the client request. In the example of FIG. 4A, it is assumed that the Main Farm Unit 204 selects subordinate Farm Unit A for servicing the client request. As shown in FIG. 2, subordinate Farm Unit A 242a has associated with it a respective server (e.g., Server A 210a) for servicing client requests which are routed to Farm Unit A 242a.

Once the Main Farm Unit has selected an appropriate subordinate Farm Unit for servicing the client request, the Main Farm Unit forwards (5) the client request to the specific server associated with the selected subordinate Farm Unit. Thus, in the example of FIG. 4A, the Main Farm Unit 204 forwards the client request to Server A 210a.

When the initial client request is received at Server A 210a, the server initiates a communication session with the client device, and generates (7) a session ID corresponding to the initiated client session. In the example of FIG. 4A, the

8

session which is initiated and associated with client 202 corresponds to an electronic commerce session initiated at the server system 230 for the client 202.

At (9), Server A creates a table entry for the initiated customer session in Data Cache A 250a and database 220. According to at least one embodiment, Data Cache A 250a may be used for storing and/or retrieving application state data and session state data relating to the customer session initiated with customer 202. The application and session state data relating to the customer session initiated for customer 202 may also be stored in the database 220. According to at least one implementation, the customer session and application state data stored on either the Data Cache A 250a or database 220 may be accessed using the session ID associated with the client 202 customer session.

FIG. 3 shows a block diagram of a specific implementation for storing customer session and application state data in data cache 250a. As shown in FIG. 3, data cache 250a may include one or more tables for storing application and/or session state data relating to selected customer sessions. In the specific implementation of FIG. 3, the data cache 250a includes a session state table 251a and one or more application state tables 253a. The session state table may be used for storing customer session state data such as, for example, customer login information. The application state tables 253a may be used for storing customer application state data such as, for example, the current contents of a customer's electronic shopping cart. Each of the plurality of data caches 250 may include data structures similar to that shown in FIG. 3. Additionally, database 220 may also include data structures which are similar to those shown in FIG. 3.

According to a specific embodiment, each of the plurality of data caches 250 is configured to store session and application state data relating to customer sessions which have been initiated with the data cache's associated server. Thus, for example, Data Cache A 250a may be configured to store and/or provide state data for customer sessions handled by Server A, and Data Cache B 250b may be configured to store and/or provide state data relating to customer sessions initiated with Server B 242b. Further, according to at least one embodiment, the database 220 may be used for storing and/or retrieving state data relating to all customer sessions which have been initiated with the host server system 230. Moreover, according to a specific implementation, the host server system 230 may be configured to store, on the database 220, the most current application and/or session state data for any given customer session.

Returning to FIG. 4A, once Server A receives the initial client request, it processes the initial client request and generates (11) an appropriate response. The processing of the client request may result in a change of the session and/or application state data associated with that client session. Accordingly, as shown in FIG. 4A, any state data which has been modified or updated by Server A will be stored (12) in the Data Cache A 250a, as well as the database 220. This procedure of writing the same data to both the data cache and the database is commonly referred to as a data write-through operation. At (13), the response generated by Server A 210a is transmitted to the Main Farm Unit 204, which then forwards (17) the response to the client 202. Before transmitting the response to the client 202, the Main Farm Unit 204 replaces (15) the source IP address of the response packet header with the IP address corresponding to the Main Farm Unit 204. One reason for changing the packet header information is that the client 202 is expecting to receive a response from Main Farm Unit 204, rather than

US 7,197,547 B1

9

from Server A 210a. As described in greater detail below, the HTTP response generated by Server A may comprise HTML data which may include one or more URLs corresponding to subordinate Farm Unit A. Additionally, as described in greater detail below, the HTTP response may also comprise cookie file data which includes the session ID corresponding to the client session, and may also include a server instance ID corresponding to the current instance of Server A.

According to at least one embodiment of the present invention, each of the plurality of servers 210 may include substantially similar content. However, each server's content may include different URLs to be provided to clients for enabling a client to access specific data from the host server system 230 via the particular server which provided the URL. For example, where the host server system 230 is configured to facilitate electronic commerce relating to on-line shopping, each of the plurality of servers 210 may include substantially similar content relating to the catalog of products which are available from the on-line merchant. The content stored on each server may differ in that each server may include one or more URLs corresponding to information which is accessible via the particular server on which the URLs reside. Thus, for example, the content on Server A 210a may include URLs for accessing content specific to Server A. Similarly, Server B 210b may include URLs for accessing content specific to Server B.

According to a specific implementation, the URLs which are included in a client response generated by a particular server correspond to an address of the specific subordinate Farm Unit (of the plurality of subordinate Farm Units 242) associated with the server which generated the client response. For example, when Server A generates a response to a client request, the response may include HTML data having at least one embedded URL. The embedded URL corresponds to an address associated with subordinate Farm Unit A 242a. As explained in greater detail below, when the client selects the embedded URL, an HTTP request is sent to subordinate Farm Unit A. As shown in the example of FIG. 2, subordinate Farm Unit A is associated exclusively with Server A, and therefore forwards the received client request to Server A. As described in greater detail below, one advantage of configuring the URLs of a particular server to correspond to an address representing that server's associated subordinate Farm Unit is that it allows the load balancing device 240 to transparently perform fail-over procedures if a failure is detected at any one of the servers 210.

According to a specific embodiment of the present invention, each server of the host server system 230 is configured to process client requests and generate appropriate responses to the requesting clients. Thus, for example, during an electronic commerce session initiated between client 202 and Server A 210a, Server A may transmit a response to client 202 comprising HTML data which includes URLs for accessing additional data from Server A. Similarly, during an electronic commerce session initiated between client 202 and Server B 210b, Server B would transmit HTML data in response to requests from client 202, wherein the HTML data includes URLs for accessing additional information from Server B. In this way, a "stickiness" scheme may be implemented in the host server system 230 whereby a specific server which is assigned (by the load balancing device 240) to respond to an initial request from a particular client also handles all subsequent requests from that client in order to make optimal use of the state data stored in the server's data cache, thereby resulting in a faster response time of the server system 230.

Returning to FIG. 4A, when the HTTP response is received at the client 202, the cookie file data (which

10

includes the session ID and Server A instance ID) is stored in a cookie file on the client machine. Additionally, the HTML data received from the host server system may be displayed to the client using, for example, a conventional web browser and display screen. At (21) the client submits a subsequent HTTP request to the host server system 230. In the specific example of FIG. 4A, it is assumed that the subsequent HTTP request is generated by the client in response to the client selecting a particular URL embedded within the HTML data displayed to the client. For example, the client may select a particular product to add to the client's electronic shopping cart. In this example, the HTTP request would correspond to an "add to cart" request to be implemented at the host server system. Since the URL selected by the client corresponds to an address of the subordinate Farm Unit A 242a, the destination of the HTTP request will be the IP address of the subordinate Farm Unit A. As shown in FIG. 4A, the HTTP request may include data relating to the cookie file stored on the client's machine such as, for example, the client session ID and Server A instance ID.

When the subordinate Farm Unit A receives the subsequent HTTP request from client 202, it forwards (23) the request to Server A. According to a specific embodiment, the load balancing device (240, FIG. 2), which includes subordinate Farm Unit A, does not perform a conventional load balancing procedure for the subsequent packet received from client 202. Since the subsequent client packet (which contains the subsequent HTTP request) was received at subordinate Farm Unit A 242a rather than the Main Farm Unit 204, the load balancing unit assumes that a client session has already been initiated between the client 202 and Server A 210a. Accordingly, the load balancing device 240 automatically forwards the subsequent client request packet to Server A, provided that a failure is not detected at Server A.

When Server A receives the subsequent HTTP request from client 202, it identifies (25) the session ID from the cookie file data transmitted along with the HTTP request. According to an alternate implementation, the subsequent client request may not include the cookie file data. When the subsequent request is received at Server A, Server A submits a request to the client to retrieve the cookie file data stored on the client machine, including, for example, the session ID and Server A instance ID. Once the cookie file data is received at Server A, the session ID corresponding the electronic commerce session for client 202 may then be identified. After identifying the session ID, Server A processes (27) the subsequent HTTP request. In processing the request, Server A may update and/or retrieve (29) state data relating to the identified session ID from the Data Cache A 250a and/or database 220. According to a specific embodiment, customer requests which do not involve a change in the customer state data may be processed by retrieving data from the server's associated data cache. For example, a customer request to display the contents of the customer's electronic shopping cart may be handled by Server A retrieving the appropriate data from the Data Cache A. It will be appreciated that the server does not need to access the database 220 in order to respond to this request. Accordingly, the processing time for responding to the client's request may be significantly reduced. However, if the processing of the client's request results in a change in the client's session and/or application state data (such as, for example, an "add to cart" request), a data write-through operation should preferably be performed, wherein the updated state data for that client is stored in both the data

US 7,197,547 B1

11

cache 250a and database 220. Once Server A has generated a response to the subsequent client request, it transmits (31) the new response to subordinate Farm Unit A. The response generated by Server A may include updated HTML data, and cookie file data which includes the session ID and Server A instance ID. When the subordinate Farm Unit A receives the response from Server A, it replaces (33) the source IP address of the packet header with the IP address of subordinate Farm Unit A. Thereafter, the subordinate Farm Unit A transmits (35) the HTTP response to the client 202.

As illustrated in the example of FIG. 4A, at (37) it is assumed that a failure occurs at Server A. According to a specific embodiment, the load balancing device 240 may be configured to detect a failure at any one of the plurality of servers 210 by implementing a Server Status Check Procedure such as that shown in FIG. 6 of the drawings.

FIG. 6 shows a flow diagram of a Server Status Check Procedure 600 in accordance with a specific embodiment of the present invention. According to a specific embodiment, a separate instance of the Server Status Check Procedure may be implemented for each server in the host server system 230, thereby allowing the load balancing device to simultaneously check the status of any desired number of servers in the server farm system. According to a specific implementation, the Server Status Check Procedure 600 may be executed at periodic intervals, or at times when the load balancing device is not busy, or may be implemented before the load balancing device forwards a received client request to the appropriate server.

In order to check the status of a selected server, the load balancing device transmits (602) to a selected server a test request (e.g., test HTTP request) which includes a test URL. According to a specific embodiment, the test URL causes the selected server to utilize desired components of the server's technology stack in order to generate a response to the test request. By analyzing and comparing the server's response against a predetermined response corresponding to a healthy server, the load balancing device is able to detect whether there exists a failure at one or more components of the selected server. Accordingly, at 604 a determination is made as to whether a server failure or server timeout has been detected. If a server failure or server timeout has been detected, the selected server is disabled (606) from receiving subsequent requests from any client. Additionally, a backup server is enabled (608) to receive any future requests directed to the disabled server. At 610 a Server Status Table may be updated to reflect the disabled status of the selected server and enabled status of the backup server. According to a specific implementation, the Server Status Table may reside at the load balancing device 240. At 618 the load balancing device may delay a random or predetermined time period before reinitiating the Server Status Check Procedure for the selected server.

Returning to block 604, if a server failure or server timeout is not detected for the selected server, at 612 a determination is made as to whether the status of the selected server reflects that it is enabled to receive client requests. If it is determined that the status of the selected server is enabled, it is assumed that the server is operating properly, and that the server may receive client request packets for processing. If, however, it is determined that the status of the selected server is disabled, it may be assumed that a failure was previously detected at the selected server, and that the selected server is now functioning properly. Accordingly, the status of the selected server is updated to enable (614) the server to receive client request packets for processing. The Server Status Table is then updated (616) to reflect the current status of the selected server.

12

FIG. 4B shows a specific embodiment of a data flow diagram corresponding to a sequence of events which may be implemented by the present invention in response to detecting a failure at a selected server of the server system 230. The example of FIG. 4B is intended to be a continuation of the example described previously with respect to FIG. 4A.

At (37), it is assumed that a failure occurs at Server A 210a. At (39) the client 202 transmits a subsequent HTTP request to subordinate Farm Unit A 242a. According to a specific implementation, subordinate Farm Unit A represents a logical device which is part of the load balancing device 240 of FIG. 2. At (41) the load balancing device 240 detects (via, for example, the Server Status Check Procedure 600 of FIG. 6) that a failure has occurred at Server A, and chooses a backup or alternate server for servicing the HTTP request received from client 202. According to a specific implementation, the selection of the alternate or backup server may be performed by utilizing a load balancing procedure or protocol which is implemented at the load balancing device 240.

According to a specific embodiment of the present invention, when a backup server takes over an initiated client session for a failed server, the initiated client session will continue at the backup server, even after the failed server subsequently recovers. After the failed server recovers, new client requests may then be directed to the recovered server, whereupon new client sessions will then be initiated.

In the example of FIG. 4B, it is assumed that the load balancing device selects Server B 210b as the alternate or backup server. Accordingly, the load balancing device 240 forwards (43) the client request packet to Server B. When Server B receives the forwarded client request, it identifies (45) the session ID from the cookie file data transmitted by client 202. Server B then attempts to access the state data associated with the electronic commerce session for client 202 from the Data Cache B 250b. However, at this point, the client 202 state information has been stored on Data Cache A 250a, but has not been stored on Data Cache B 250b. Accordingly, Server B will detect (49) a cache miss when attempting to retrieve data relating to the identified session ID from the Data Cache B. In response to detecting a cache miss, each of the plurality of servers 210 may be configured to retrieve session and application state data relating to an identified session ID from the database 220 into its local data cache. Thus, as shown in FIG. 4B, Server B retrieves (51) the session and application state data relating to the client 202 session ID from the database 220, and stores (53) the retrieved session and application state data in the Data Cache B. Thereafter, Server B processes (55) the client request, and stores (57) any updated session and/or application state data (relating to that session ID) in the appropriate state table(s) of the Data Cache B 250b and database 220.

At (59), Server B transmits its response to the client request to the subordinate Farm Unit A, whereupon the subordinate Farm Unit A replaces (61) the source IP address of the packet header with the IP address of subordinate Farm Unit A, and then transmits (63) the HTTP response (generated by Server B) to the client 202.

As stated previously, each response generated by a particular server may include one or more URLs corresponding to an address of the particular subordinate Farm Unit associated with that server. Thus, the response generated by Server B may include at least one URL corresponding to an address associated with subordinate Farm Unit B.

US 7,197,547 B1

13

At (65), it is assumed that the client selects one of the URLs provided in the response generated by Server B. The selection of this URL causes the client machine to transmit an appropriate request to subordinate Farm Unit B, whereupon it is then forwarded to Server B for processing in a manner similar to events 23–35, described previously with respect to FIG. 4A.

According to a specific embodiment, Server B will continue to receive and process subsequent request packets from client 202, even after Server A has recovered from its failure and come back on-line. Alternatively, when Server A recovers from its failure, and is detected as functioning normally, Server A may then be enabled to receive subsequent request packets from client 202. However, in this latter situation it is possible for the Data Cache A 250a to have old or erroneous data relating to the current status and/or state of the client 202 electronic commerce session. Accordingly, in order to ensure that Server A uses the most up-to-date state information relating to the client 202 session, a Server Instance ID Verification Procedure may be initiated as described, for example, in FIG. 5 of the drawings.

FIG. 5 shows a flow diagram of a Server Instance ID Verification Procedure 500 in accordance with a specific embodiment of the present invention. According to a specific embodiment, a separate instance of the Server Instance ID Verification Procedure may be implemented on each of the plurality of servers 210. Further each instance of the Server Instance ID Verification Procedure may be executed at its respective server each time a client request is received at that server.

According to a specific embodiment, each of the plurality of servers 210 has associated with it a unique server instance ID representing the current instance of that particular server session. Each time a server is rebooted or recovers from a failure, the server instance ID associated with that server changes. The current server instance ID may be stored locally at the server, and may also be included in each response generated by that server in response to a client request. The server instance ID transmitted to the client may be stored in the cookie file of the client machine. When the client submits a request to the host server system 230, the request may include cookie file data such as, for example, the server instance ID.

Referring to FIG. 5, at 502 a request is received from a particular client at a specific server of the server farm system. The server identifies (504) the server instance ID from the client cookie data transmitted from the client machine. Additionally, the server retrieves (506) the current server instance ID, which may be stored on the server. A determination is then made (508) as to whether the cookie server instance ID (provided by the client machine) is the same as the current server instance ID (provided by the server). If both server instance IDs are identical, then it may be assumed that the server has not experienced a failure or been rebooted since last communicating with the client, and that the client state data stored in the server's data cache is current and up-to-date. Accordingly, the client request will be processed (512) normally.

However, if it is determined that the cookie server instance ID is not the same as the current server instance ID of the selected server, then it may be assumed that the server has either experienced a failure or been rebooted since last communicating with the requesting client. Accordingly, it is likely that the data in the server's data cache relating to the client's electronic commerce session is not up-to-date. In response, the client state data stored on the server's data

14

cache (which may be identified using the client session ID) is flushed (510). Thereafter, the client request is processed (512) normally as described, for example, with respect to FIG. 4C of the drawings.

FIG. 4C shows a specific embodiment of a data flow diagram which illustrates a sequence of events that may occur after a particular server has been rebooted or has recovered from a failure. The example of FIG. 4C is intended to be a continuation of the example described previously with respect to FIG. 4A.

As shown in FIG. 4C, at (71) it is assumed that Server A recovers from a failure, and is reinitiated with a new server instance ID. At (73), the load balancing device 240 detects that Server A has recovered from its failure, and changes the status of Server A to enable it to again receive client requests. At (75) client 202 submits an HTTP request to subordinate Farm Unit A. Upon receiving the client request, the subordinate Farm Unit A forwards (79) the request to Server A. Server A identifies (80) the session ID and server instance ID from the client cookie file data provided by the client machine. At (81), Server A indicates the Server Instance ID Verification Procedure 500 (FIG. 5), which results in the flushing of data of Data Cache A associated with the identified session ID. At (83), Server A attempts to access the client state data for the identified session ID from Data Cache A, and detects (85) a cache miss. The session and application state data relating to the client session ID is then retrieved (86) from the database 220, and stored (87) in the Data Cache A. Server A then processes (88) the client request, and stores (89) any updated state data relating to the client session in the Cache A 250a and database 220.

At (90), Server A transmits its response to subordinate Farm Unit A, which replaces (91) the source IP address of the response packet header with the IP address of subordinate Farm Unit A, and then transmits (92) the response to the client 202.

Other Embodiments

Generally, the load balanced server farm system of the present invention may be implemented via software and/or hardware. In a specific embodiment of this invention, the technique of the present invention may be implemented in software such as an operating system or in an application running on an operating system.

A software or software/hardware hybrid load balanced server farm system of this invention may be implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. Such programmable machine may be a network device designed to handle network traffic. Such network devices typically have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, the load balanced server farm system may be implemented on one or more general-purpose network host machines such as a personal computers or workstations. Further, the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to FIG. 7, a server device 700 suitable for implementing a server of present invention includes a master central processing unit (CPU) 702, one or more interfaces 708, and a bus 707 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 702 is responsible for such tasks as processing HTTP requests, dynamically generating HTML data, generating updated

US 7,197,547 B1

15

session and application state data, accessing data from a data cache and/or persistent memory, etc. It preferably accomplishes all these functions under the control of software including an operating system and any appropriate applications software. CPU 702 may include one or more processors 703 such as a processor from the Motorola family of microprocessors or the Intel family of microprocessors. In an alternative embodiment, processor 703 is specialty designed hardware for controlling the operations of the server device 700. In a specific embodiment, a memory 701 (such as non-volatile RAM and/or ROM) also forms part of CPU 702. However, there are many different ways in which memory could be coupled to the system. Memory block 701 may be used for a variety of purposes such as, for example, caching and/or storing client session and application state data, programming instructions, etc.

The load balanced server farm system of the present invention may also employ one or more memories or memory modules (such as, for example, memory block 705) configured to store various data, program instructions, etc. The program instructions may control the operation of an operating system and/or one or more applications. The memory or memories may also be configured to store the various types of data described in this application, such as for example, HTML data, client session and application state data, etc.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

Although several preferred embodiments of this invention have been described in detail herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to these precise embodiments, and at various changes and modifications may be effected therein by one skilled in the art without departing from the scope of spirit of the invention as defined in the appended claims.

What is claimed is:

1. A load balanced server farm system for effecting electronic commerce over a data network, the system comprising:

- a load balancing system; and
- a plurality of servers in communication with the load balancing system, wherein each server includes a respective data cache;

wherein the load balancing system is configured or designed to select, using a load balancing protocol, an available first server of the plurality of servers to process an initial packet received from a source device, and is further configured or designed to route subsequent packets received from the source device to the first server,

wherein the first server includes a first data cache and is configured or designed to use state data stored on the

16

first data cache for generating a first response to the source device, the state data corresponding to an electronic commerce session initiated between the source device and the first server,

wherein the first server is configured or designed to initiate an electronic commerce session for the source device, and is configured or designed to generate session ID information including a session ID associated with the session initiated for the source device,

wherein the first server is configured or designed to include the session ID information in the first response transmitted to the source device,

wherein the first server is configured or designed to detect a cache miss in response to an attempt to access state data at the first data cache, said state data relating to the electronic commerce session associated with the session ID, and

wherein the first server is configured or designed to retrieve state data associated with the session ID from a database into the first data cache in response to detecting said cache miss.

2. The system of claim 1:

wherein the load balancing system includes a main load balancing device and a plurality of subordinate load balancing devices, each of the subordinate load balancing devices being associated with a different server farm of the server farm system; and

wherein each server of the plurality of servers is associated with a respective server farm system;

said plurality of server farms comprising a first server farm which includes a first subordinate load balancing device and the first server;

said plurality of server farms further comprising a second server farm which includes a second subordinate load balancing device and a second server, the second server including a second data cache.

3. The system of claim 2 further comprising a persistent memory database, accessible by each of the plurality of servers, for storing state information relating to at least one electronic commerce session initiated between a customer and the server farm system.

4. The system of claim 1 wherein each of the plurality of data caches includes memory for storing session and application state data relating to the initiated electronic commerce session.

5. The system of claim 1 wherein the first server is further configured or designed to store, in a persistent memory database, state data relating to the electronic commerce session initiated for the source device.

6. The system of claim 1 wherein said state data includes session state data and application state data corresponding to the electronic commerce session for the source device.

7. The system of claim 1 wherein at least one of the subsequent packets received from the source device includes the session ID corresponding to the initiated electronic commerce session; and

wherein the first server is further configured or designed to access said state data from said data cache using said session ID.

8. The system of claim 1:

wherein at least one of the subsequent packets received from the source device includes a session ID corresponding to an electronic commerce session initiated at the server farm system, said electronic commerce session being associated with the source device;

US 7,197,547 B1

17

wherein said first server is further configured or designed to access session state information from the first data cache using the session ID, the session state information relating to the electronic commerce session associated with the source device; and

wherein said first server is further configured or designed to generate a second response to the subsequent request packet using said retrieved state information, wherein said second response includes updated state information relating to said electronic commerce session.

9. The system of claim 8 wherein said first server is further configured or designed to store the updated state information generated by the first server in the first data cache, and is further configured or designed to store the updated state information generated by the first server in a persistent memory database.

10. The system of claim 1 wherein the first response includes information relating to an instance ID associated with the first server.

11. The system of claim 10:

wherein at least one subsequent packet received from the source device includes the server instance ID information;

wherein said first server is further configured or designed to compare the server instance ID information received from the source device with server instance ID information stored at the first server; and

wherein said first server is further configured or designed to flush at least a portion of the first data cache if the server instance ID provided by the source device is not equivalent to the server instance ID stored at the first server.

12. The system of claim 1 wherein said first server is further configured or designed to store the state data retrieved from the database in the first data cache.

13. The system of claim 1 wherein each of the plurality of servers includes substantially similar content, and wherein the content on each server includes a respective plurality of Uniform Resource Locators (URLs) for accessing information specific to that server.

14. The system of claim 1 wherein said first server is further configured or designed to provide at least one URL in said first response, wherein URL corresponds to an address for accessing information via the first server.

15. The system of claim 2 wherein said first server is further configured or designed to provide first server data to at least one requesting client, the first server data including at least one URL corresponding to at least one address for accessing information via the first server; and

wherein said second server is further configured or designed to provide second server data to at least one requesting client, the second server data including at least one URL corresponding to at least one address for accessing information via the second server.

16. The system of claim 2 wherein said initial request is received at the main load balancing device; and

wherein at least a portion of said subsequent packet are received at the first subordinate load balancing device.

17. The system of claim 1 wherein the load balancing system is further configured or designed to route said subsequent request packets directly to the first server without performing said load balancing procedure.

18. A method for effecting electronic commerce over a data network, the method comprising:

receiving an initial request packet from a source device; performing a load balancing protocol, wherein the initial packet from source device is assigned to a first server

18

of a load balanced server system, the load balanced server system including a plurality of servers, the first server including a first data cache;

generating a first response to the initial request packet at the first server;

transmitting the first response to the source device; and causing subsequent packets received from said source device to be routed to the first server,

wherein the method further comprises:

initiating a communication session between the source device and the first server;

generating session ID information associated with the session initiated with the source device;

including the session ID information in the first response transmitted to the source device;

storing, in the first data cache, state data relating to the session initiated between the source device and the first server;

detecting a cache miss in response to an attempt to access state data on the first data cache, said state data relating to the session associated with the ID information; and

retrieving state data associated with the session ID information from a database into the first data cache in response to detecting said cache miss.

19. The method of claim 18 wherein said causing includes:

receiving subsequent request packets from the source device; and

routing said subsequent request packets to the first server.

20. The method of claim 18 further comprising storing, in a persistent memory database, state data relating to the session initiated between the source device and the first server.

21. The method of claim 18 wherein said state data includes session state data and application state data corresponding to the communication session between the source device and the first server.

22. The method of claim 18 further comprising accessing said state data using said session ID information.

23. The method of claim 18 further comprising:

receiving a second request packet from the source device, the second packet including session ID information for identifying the initiated session between the source device and the first server;

routing said second request packet to the first server; accessing session state information from the first data cache using the session ID information, the session state information relating to the session between the source device and the first server;

generating a second response to the second request packet using said retrieved state information, wherein said second response includes updated state information relating to said source device session; and

transmitting said second response the source device.

24. The method of claim 23 further comprising:

storing the updated state information generated by the first server in the first data cache; and

storing the updated state information generated by the first server in a persistent memory database.

25. The method of claim 18 wherein the first response includes information relating to an instance ID associated with the first server.

26. The method of claim 25 further comprising:

receiving a second request packet from the source device, the second packet including said server instance ID information;

US 7,197,547 B1

19

comparing the server instance ID information received from the source device with server instance ID information stored at the first server; and

flushing at least a portion of the first data cache if the server instance ID provided by the source device is not equivalent to the server instance ID stored at the first server.

27. The method of claim 18 wherein said retrieving comprises storing the state data retrieved from the database in the first data cache.

28. The method of claim 18 wherein each of the plurality of servers includes substantially similar content, and wherein the content on each server includes a respective plurality of Uniform Resource Locators (URLs) for accessing information specific to that server.

29. The method of claim 18 further comprising providing at least one URL in said first response, wherein the at least one URL corresponds to at least one address for accessing information via the first server.

30. The method of claim 18 wherein said initial request is received at a main farm device which performs said load balancing protocol; and

wherein said subsequent packet receiving includes receiving a second request packet from said source device at a sub-farm device associated with the first server.

31. The method of claim 18:

wherein said initial request is received at a main logical load balancing unit; and

wherein the subsequent request packets are received at a subordinate logical load balancing unit associated with the first server.

32. The method of claim 18 further comprising routing said subsequent request packets directly to the first server without performing said load balancing protocol.

33. A computer program product comprising a computer readable medium, the computer readable medium including computer code for implementing the method of claim 18.

34. A system for effecting electronic commerce over a data network, the system comprising:

means for receiving an initial request packet from a source device;

means for performing a load balancing protocol, wherein the initial packet from source device is assigned to a first server of a load balanced server system, the load balanced server system including a plurality of servers, the first server including a first data cache;

20

means for generating a first response to the initial request packet at the first server;

means for transmitting the first response to the source device; and

means for causing subsequent packets received from said source device to be routed to the first server;

wherein the system further comprises:

means for initiating a communication session between the source device and the first server;

means for generating session ID information associated with the session initiated with the source device;

means for including the session ID information in the first response transmitted to the source device;

means for storing, in the first data cache, state data relating to the session initiated between the source device and the first server;

means for detecting a cache miss in response to an attempt to access state data on the first data cache, said state data relating to the session associated with the ID information; and

means for retrieving state data associated with the session ID information from a database into the first data cache in response to detecting said cache miss.

35. The system of claim 34 further comprising:

means for receiving a second request packet from the source device, the second packet including session ID information for identifying the initiated session between the source device and the first server;

means for routing said second request packet to the first server;

means for accessing session state information from the first data cache using the session ID information, the session state information relating to the session between the source device and the first server;

means for generating a second response to the second request packet using said retrieved state information, wherein said second response includes updated state information relating to said source device session; and

means for transmitting said second response the source device.

* * * * *