



US006889181B2

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

(10) **Patent No.:** **US 6,889,181 B2**
(45) **Date of Patent:** ***May 3, 2005**

(54) **NETWORK FLOW SWITCHING AND FLOW DATA EXPORT**

4,131,767 A 12/1978 Weinstein
4,161,719 A 7/1979 Parikh et al.
4,316,284 A 2/1982 Howson

(75) Inventors: **Darren R. Kerr**, Union City, CA (US);
Barry L. Bruins, Los Altos, CA (US)

(Continued)

(73) Assignee: **Cisco Technology, Inc.**, San Jose, CA (US)

FOREIGN PATENT DOCUMENTS

(*) Notice: This patent issued on a continued prosecution application filed under 37 CFR 1.53(d), and is subject to the twenty year patent term provisions of 35 U.S.C. 154(a)(2).

EP 0 384 758 A2 8/1990
EP 0 431 751 A1 6/1991
EP 0 567 217 A2 10/1993
WO WO93/07569 4/1993
WO WO93/07692 4/1993
WO WO94/01828 1/1994
WO WO95/20850 8/1995

OTHER PUBLICATIONS

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

Awerbuch et al., "Sparse Partitions," IEEE, 1990, pp. 503-513.*

Bouloutas et al., "Some Graph Partitioning Problems and Algorithms Related to Raiting in Large Computer Networks," IEEE, 1989, pp. 362-370.*

(21) Appl. No.: **08/886,900**

(Continued)

(22) Filed: **Jul. 2, 1997**

(65) **Prior Publication Data**

US 2001/0051864 A1 Dec. 13, 2001

Primary Examiner—Thai Phan

(74) *Attorney, Agent, or Firm*—Hickman Palermo Truong & Becker LLP

Related U.S. Application Data

(63) Continuation of application No. 08/771,438, filed on Dec. 20, 1996, and a continuation of application No. 08/655,429, filed on May 28, 1996, now Pat. No. 6,243,667.

(51) **Int. Cl.**⁷ **G06F 9/34**

(52) **U.S. Cl.** **703/27; 703/20; 370/235; 370/389; 370/392; 370/353**

(58) **Field of Search** 395/500, 858, 395/200.58; 370/351, 355, 229, 235, 238, 379, 389, 392, 353; 340/825.04; 703/20, 27

ABSTRACT


The invention provides a method and system for switching in networks responsive to message flow patterns. A message "flow" is defined to comprise a set of packets to be transmitted between a particular source and a particular destination. When routers in a network identify a new message flow, they determine the proper processing for packets in that message flow and cache that information for that message flow. Thereafter, when routers in a network identify a packet which is part of that message flow, they process that packet according to the proper processing for packets in that message flow. The proper processing may include a determination of a destination port for routing those packets and a determination of whether access control permits routing those packets to their indicated destination.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,064,360 A 12/1977 Koenig

22 Claims, 3 Drawing Sheets



U.S. PATENT DOCUMENTS

4,397,020 A	8/1983	Howson	5,199,049 A	3/1993	Wilson
4,419,728 A	12/1983	Larson	5,206,638 A	4/1993	McKeown
4,424,565 A	1/1984	Larson	5,206,886 A	4/1993	Bingham
4,437,087 A	3/1984	Petr	5,208,811 A	5/1993	Kashio et al.
4,438,511 A	3/1984	Baran	5,212,686 A	5/1993	Joy et al.
4,439,763 A	3/1984	Limb	5,224,099 A	6/1993	Corbalis et al.
4,445,213 A	4/1984	Baugh et al.	5,226,120 A	7/1993	Brown et al.
4,446,555 A	5/1984	Devault et al.	5,228,062 A	7/1993	Bingham
4,456,957 A	6/1984	Schieltz	5,229,994 A	7/1993	Balzano et al.
4,464,658 A	8/1984	Thelen	5,231,633 A	7/1993	Hluchyj et al.
4,491,945 A	1/1985	Turner	5,233,604 A	8/1993	Ahmadi et al.
4,499,576 A	2/1985	Fraser	5,237,564 A	8/1993	Lespagnol et al.
4,506,358 A	3/1985	Montgomery	5,241,682 A	8/1993	Bryant et al.
4,507,760 A	3/1985	Fraser	5,243,342 A	9/1993	Kattemalalavadi et al.
4,532,626 A	7/1985	Flores et al.	5,243,596 A	9/1993	Port et al.
4,644,532 A	2/1987	George et al.	5,247,516 A	9/1993	Bernstein et al.
4,646,287 A	2/1987	Larson et al.	5,249,178 A	9/1993	Kurano et al.
4,677,423 A	6/1987	Benvenuto et al.	5,249,292 A	9/1993	Chiappa
4,679,189 A	* 7/1987	Olson et al. 370/60	5,251,205 A	10/1993	Callon
4,679,227 A	7/1987	Hughes-Hartogs	5,253,251 A	10/1993	Aramaki
4,713,806 A	12/1987	Oberlander et al.	5,255,291 A	10/1993	Holden et al.
4,723,267 A	2/1988	Jones et al.	5,260,933 A	11/1993	Rouse
4,731,816 A	3/1988	Hughes-Hartogs	5,260,978 A	11/1993	Fleischer et al.
4,750,136 A	6/1988	Arpin et al.	5,268,592 A	12/1993	Bellamy et al.
4,757,495 A	7/1988	Decker et al.	5,268,900 A	12/1993	Hluchyj et al.
4,763,191 A	8/1988	Gordon et al.	5,271,004 A	12/1993	Proctor et al.
4,769,810 A	9/1988	Eckberg, Jr. et al.	5,274,631 A	12/1993	Bhardwaj
4,769,811 A	9/1988	Eckberg, Jr. et al.	5,274,635 A	12/1993	Rahman et al.
4,771,425 A	9/1988	Baran et al.	5,274,643 A	12/1993	Fisk
4,788,657 A	11/1988	Douglas et al.	5,280,470 A	1/1994	Buhrke et al.
4,819,228 A	4/1989	Baran et al.	5,280,480 A	1/1994	Pitt et al.
4,827,411 A	5/1989	Arrowood et al.	5,280,500 A	1/1994	Mazzola et al.
4,833,706 A	5/1989	Hughes-Hartogs	5,283,783 A	2/1994	Nguyen et al.
4,835,737 A	5/1989	Herrig et al.	5,287,103 A	2/1994	Kasprzyk et al.
4,879,551 A	11/1989	Georgiou et al.	5,287,453 A	* 2/1994	Roberts 395/200
4,893,304 A	1/1990	Giacopelli et al.	5,291,482 A	3/1994	McHarg et al.
4,893,306 A	1/1990	Chao et al.	5,301,333 A	4/1994	Lee
4,903,261 A	2/1990	Baran et al.	5,305,311 A	4/1994	Lyles
4,905,233 A	2/1990	Cain et al.	5,307,343 A	4/1994	Bostica et al.
4,922,486 A	5/1990	Lidinsky et al.	5,309,437 A	5/1994	Perlman et al. 730/85.13
4,933,937 A	6/1990	Konishi	5,311,509 A	5/1994	Heddes et al.
4,960,310 A	10/1990	Cushing	5,313,454 A	5/1994	Bustini et al.
4,962,497 A	10/1990	Ferenc et al.	5,313,582 A	5/1994	Hendel et al.
4,962,532 A	10/1990	Kasirai et al.	5,317,562 A	5/1994	Nardin et al.
4,965,772 A	10/1990	Daniel et al.	5,319,644 A	6/1994	Liang
4,970,678 A	11/1990	Sladowski et al.	5,325,358 A	6/1994	Goeldner
4,979,118 A	12/1990	Kheradpir 364/436	5,327,421 A	7/1994	Hiller et al.
4,980,897 A	12/1990	Decker et al.	5,331,637 A	7/1994	Francis et al.
4,991,169 A	2/1991	Davis et al.	5,335,224 A	8/1994	Cole et al.
5,003,595 A	3/1991	Collins et al.	5,339,311 A	8/1994	Turner
5,014,265 A	5/1991	Hahne et al.	5,345,445 A	9/1994	Hiller et al.
5,014,327 A	5/1991	Potter et al.	5,345,446 A	9/1994	Hiller et al.
5,020,058 A	5/1991	Holden et al.	5,353,283 A	10/1994	Tsuchiya
5,033,076 A	7/1991	Jones et al.	5,357,506 A	10/1994	Sugawara
5,051,987 A	9/1991	Conlon	5,359,592 A	10/1994	Corbalis et al.
5,054,034 A	10/1991	Hughes-Hartogs	5,361,250 A	11/1994	Nguyen et al.
5,059,925 A	10/1991	Weisbloom	5,361,256 A	11/1994	Doeringer et al.
5,063,612 A	11/1991	McKeown	5,361,259 A	11/1994	Hunt et al.
5,072,449 A	12/1991	Enns et al.	5,365,524 A	11/1994	Hiller et al.
5,088,032 A	2/1992	Bosack	5,367,517 A	11/1994	Cidon et al.
5,095,480 A	* 3/1992	Fenner 370/94.1	5,371,852 A	12/1994	Attanasio et al.
5,101,402 A	3/1992	Chiu et al.	5,377,181 A	12/1994	Rogers
RE33,900 E	4/1992	Howson	5,381,413 A	1/1995	Tobagi et al.
5,115,431 A	5/1992	Williams et al.	5,386,413 A	1/1995	McAuley et al.
5,115,495 A	5/1992	Tsuchiya et al.	5,386,567 A	1/1995	Lien et al.
5,128,926 A	7/1992	Perlman et al.	5,390,170 A	2/1995	Sawant et al.
5,128,945 A	7/1992	Enns et al.	5,390,175 A	2/1995	Hiller et al.
5,136,580 A	8/1992	Vidlock et al.	5,392,280 A	2/1995	Zheng
5,166,930 A	11/1992	Braff et al.	5,394,394 A	2/1995	Crowther et al.
5,189,662 A	2/1993	Kleine-Altekamp	5,394,402 A	2/1995	Ross
			5,400,325 A	3/1995	Chatwani et al.

US 6,889,181 B2

Page 3

5,408,469 A	4/1995	Opher et al.	5,634,011 A	5/1997	Auerbach et al.
5,416,842 A	5/1995	Aziz	5,638,359 A	6/1997	Peltola et al.
5,422,880 A	6/1995	Heitkamp et al.	5,644,713 A	7/1997	Makishima
5,422,882 A	6/1995	Hiller et al.	5,644,718 A	7/1997	Belove et al.
5,423,002 A	6/1995	Hart	5,651,002 A	7/1997	Van Seters et al.
5,426,636 A	6/1995	Hiller et al.	5,659,684 A	8/1997	Giovannoni et al.
5,428,607 A	6/1995	Hiller et al.	5,666,353 A	9/1997	Klausmeier et al.
5,430,715 A	7/1995	Corbalis et al.	5,673,265 A	9/1997	Gupta et al.
5,430,729 A	7/1995	Rahnema	5,675,579 A	10/1997	Watson et al.
5,432,784 A	7/1995	Ozveren	5,678,006 A	10/1997	Valizadeh et al.
5,434,863 A	7/1995	Onishi et al.	5,680,116 A	10/1997	Hashimoto et al.
5,436,886 A	7/1995	McGill	5,684,797 A	11/1997	Aznar et al.
5,440,550 A	8/1995	Follett	5,684,954 A	11/1997	Kaiserswerth et al.
5,442,457 A	8/1995	Najafi	5,687,324 A	11/1997	Green et al.
5,442,624 A	8/1995	Bonomi et al.	5,689,506 A	11/1997	Chiussi et al.
5,442,630 A	8/1995	Gagliardi et al.	5,691,984 A	11/1997	Gardner et al.
5,448,559 A	9/1995	Hayter et al.	5,694,390 A	12/1997	Yamato et al.
5,452,294 A	9/1995	Natarajan	5,699,532 A	12/1997	Barrett
5,452,297 A	9/1995	Hiller et al.	5,724,351 A	3/1998	Chao et al.
5,473,599 A	12/1995	Li et al.	5,740,157 A	4/1998	Demiray
5,473,607 A	12/1995	Hausman et al.	5,742,760 A	4/1998	Picazo, Jr. et al.
5,477,541 A	12/1995	White et al.	5,742,905 A	4/1998	Pepe et al.
5,485,455 A	1/1996	Dobbins et al.	5,748,186 A *	5/1998	Raman 345/302
5,490,140 A	2/1996	Abensour et al.	5,748,617 A	5/1998	McLain, Jr.
5,490,258 A *	2/1996	Fenner 395/401	5,754,547 A	5/1998	Nakazawa
5,491,687 A	2/1996	Christensen et al.	5,771,231 A	6/1998	Watanabe
5,491,804 A	2/1996	Heath et al.	5,774,698 A	6/1998	Olnowich
5,497,368 A	3/1996	Reijnierse et al.	5,781,772 A	7/1998	Wilkinson, III
5,497,371 A	3/1996	Ellis et al.	5,793,745 A	8/1998	Manchester
5,500,858 A	3/1996	McKeown	5,793,978 A	8/1998	Fowler
5,504,747 A	4/1996	Sweasey	5,796,736 A	8/1998	Suzuki
5,509,006 A	4/1996	Wilford et al.	5,802,054 A	9/1998	Bellenger
5,509,123 A	4/1996	Dobbins et al.	5,809,415 A	9/1998	Rossmann
5,517,494 A	5/1996	Green	5,835,696 A	11/1998	Hess
5,517,662 A	5/1996	Hashimoto et al.	5,835,710 A	11/1998	Nagami et al.
5,519,704 A	5/1996	Farinacci et al.	5,842,224 A	11/1998	Fenner
5,519,858 A	5/1996	Walton et al. 395/600	5,854,903 A	12/1998	Morrison et al.
5,526,489 A	6/1996	Nilakantan et al.	5,856,981 A	1/1999	Voelker
5,530,963 A	6/1996	Moore et al.	5,859,837 A	1/1999	Crayford
5,535,195 A	7/1996	Lee	5,872,783 A	2/1999	Chinetal
5,539,734 A	7/1996	Burwell et al.	5,892,924 A	4/1999	Lyon et al. 395/200.75
5,539,737 A	7/1996	Lo et al.	5,898,686 A	4/1999	Virgile
5,539,743 A	7/1996	Burwell et al.	5,903,559 A	5/1999	Acharya et al.
5,541,911 A	7/1996	Nilakantan et al.	5,905,723 A	5/1999	Varghese et al.
5,546,370 A	8/1996	Ishikawa	5,909,440 A	6/1999	Ferguson
5,548,593 A	8/1996	Peschi	5,914,953 A	6/1999	Krause et al.
5,550,816 A	8/1996	Hardwick et al.	5,917,820 A	6/1999	Rekhter
5,555,244 A	9/1996	Gupta et al.	5,917,821 A	6/1999	Gobuyan et al.
5,559,801 A	9/1996	Lo	5,920,566 A	7/1999	Hendel et al.
5,561,669 A	10/1996	Lenney et al.	5,970,232 A	10/1999	Passint et al.
5,566,170 A	10/1996	Bakke et al.	5,991,817 A	11/1999	Rowett et al.
5,572,522 A	11/1996	Calamvokis et al.	5,996,021 A	11/1999	Civanlar
5,577,035 A	11/1996	Hayter et al.	6,006,264 A	12/1999	Colby et al.
5,583,862 A	12/1996	Callon	6,011,795 A	1/2000	Varghese et al.
5,586,121 A	12/1996	Moura et al.	6,014,659 A	1/2000	Wilkinson, III
5,592,470 A	1/1997	Rudrapatna et al.	6,018,524 A	1/2000	Turner et al.
5,598,115 A	1/1997	Holst	6,023,733 A	2/2000	Periasamy et al.
5,598,532 A	1/1997	Liron	6,047,323 A *	4/2000	Krause 709/227
5,598,581 A	1/1997	Daines et al.	6,052,683 A	4/2000	Irwin
5,600,798 A	2/1997	Cherukuri et al.	6,055,364 A	4/2000	Speakman
5,604,868 A	2/1997	Komine et al.	6,061,712 A	5/2000	Tzeng
5,608,726 A	3/1997	Virgile	6,067,574 A	5/2000	Tzeng
5,613,096 A	3/1997	Danknick	6,078,590 A	6/2000	Farinacci et al.
5,617,417 A	4/1997	Sathe et al.	6,101,115 A	8/2000	Ross
5,617,421 A	4/1997	Chin et al.	6,111,877 A	8/2000	Wilford
5,625,622 A	4/1997	Johri	6,122,279 A	9/2000	Milway et al.
5,630,125 A *	5/1997	Zellweger 395/614	6,147,996 A	11/2000	Laor
5,631,908 A	5/1997	Saxe	6,167,438 A *	12/2000	Yates et al. 709/216
5,632,021 A	5/1997	Jennings et al.	,021,295 A1	8/2001	Wilford
5,633,869 A	5/1997	Burnett et al.	6,157,641 A1	12/2001	Wilford
5,634,010 A	5/1997	Ciscon et al.			

OTHER PUBLICATIONS

- Chandranmenon and Varghese, "Trading Packet Headers for Packet Processing," IEEE, Apr. 96, pp. 141-52.*
- Allen, M., "Novell IPX Over Various WAN Media (IPXWAN)," Network Working Group, RFC 1551, Dec. 1993, pp. 1-22.
- Becker, D., "3c589.c: A 3c589 EtherLink3 ethernet driver for linux," becker @CESDIS.gsfc.nasa.gov, May 3, 1994, pp. 1-13.
- Chowdhury, et al., "Alternative Bandwidth Allocation Algorithms for Packet Video in ATM Networks," INFOCOM 1992, pp. 1061-1068.
- Doeringer, W., "Routing on Longest-Matching Prefixes," IEEE/ACM Transactions in Networking, vol. 4, No. 1, Feb. 1996, pp. 86-97.
- Esaki, et al., "Datagram Delivery in an ATM-Internet," 2334b IEICE Transactions on Communications, Mar. 1994, No. 3, Tokyo, Japan.
- IBM Corporation, "Method and Apparatus for the Statistical Multiplexing of Voice, Data and Image Signals," IBM Technical Disclosure Bulletin, No. 6, Nov. 1992, pp. 409-411.
- Pei, et al., "Putting Routing Tables in Silicon," IEEE Network Magazine, Jan. 1992, pp. 42-50.
- Perkins, D., "Requirements for an Internet Standard Point-to-Point Protocol," Network Working Group, RFC 1547, Dec. 1993, pp. 1-19.
- Simpson, W., "The Point-to-Point Protocol (PPP)," Network Working Group, RFC 1548, Dec. 1993, pp. 1-53.
- Tsuchiya, P.F., "A Search Algorithm for Table Entries with Non-Contiguous Wildcarding," Abstract, Bellcore.
- Zhang, et al., "Rate-Controlled Static-Priority Queueing," INFOCOM 1993, pp. 227-236.
- Williams Stallings, Data and Computer Communications, pp. 329-333, Prentice Hall, Upper Saddle River, New Jersey 07458.
- Anderson, Thomas E. et al. "High Speed Switch Scheduling for Local-Area networks". IEEE. pp. 319-352. Dated Nov. 1993.
- McKeown, Nicholas et al. "Sceduling Cells in an Input-Queued Switch". IEEE Electronics Letters. Dec. 9, 1993, pp. 2174-5.
- Newman et al. "Flow Fabelled IP: A Connectionless Approach to ATM". IEEE, 1996, pp. 12 51-60.
- Prabhaker, Balaji et al. "Multicast Scheduling for Input-Queued Switches". IEEE Jounal of Selected Areas in Communications, vol. 15, No. 5, pages: 1-23, 855-866.
- Le, My T. et al. The BayBridge: A High Speed Bridge/Router Between FDDI and SMDS part II-the SMDS interface and DQDB MAC Chip. IEEE Trans Networking (to appear).
- McKeown, Nicholas et al. "The BayBridge: A High Speed brdige/Router Between FDDI and SMDS part I-Architecture and Performance". IEEE Trans Networking (to appear).
- Karol, Mark J. et al. "Input Output Queuing on a Space-Division Packet Switch". IEEE. pp. 59-68. Dated Dec. 1987.
- Chuang, Shang-Tse et al. Matching Output Queuing with a Combined Input and Output Queued Switch pp. 1-25.
- Demers, A. et al. "Analysis and Simulation of a Fair Queuing Algorithm". Proceedings ACM SIGcomm, pp. 1-12, Sep. 1989.
- Duffy, JIM. "Cisco Goes With The Flow". Network World, Apr. 22, 1996.
- Girish et al. "Trading Packet Headers for Packet Processing". IEEE, 1996, pp. 141-152.

* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

1

NETWORK FLOW SWITCHING AND FLOW DATA EXPORT

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of the following applications:

application Ser. No. 08/655,429, filed May 28, 1996, in the name of the same inventors and having the same assignee, titled "Network Flow Switching and Flow Data Export"; and

application Ser. No. 08/771,438, filed Dec. 20, 1996, in the name of the same inventors and having the same assignee, titled "Network Flow Data Export".

Both of these applications are hereby incorporated by reference as if fully set forth herein.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to network switching and data export responsive to message flow patterns.

2. Description of Related Art

In computer networks, it commonly occurs that message traffic between a particular source and a particular destination will continue for a time with unchanged routing or switching parameters. For example, when using the file-transfer protocol "FTP" there is substantial message traffic between the file's source location and the file's destination location, comprising the transfer of many packets which have similar headers, differing in the actual data which is transmitted. During the time when message traffic continues, routing and switching devices receiving packets comprising that message traffic must examine those packets and determine the processing thereof.

One problem which has arisen in the art is that processing demands on routing and switching devices continue to grow with increased network demand. It continues to be advantageous to provide techniques for processing packets more quickly. This problem has been exacerbated by addition of more complex forms of processing, such as the use of access control lists.

It would therefore be advantageous to provide techniques in which the amount of processing required for any individual packet could be reduced. With inventive techniques described herein, information about message flow patterns is used to identify packets for which processing has already been determined, and therefore to process those packets without having to re-determine the same processing. The amount of processing required for any individual packet is therefore reduced.

Information about message flow patterns would also be valuable for providing information about use of the network, and could be used for a variety of purposes by network administrators, routing devices, service providers, and users.

Accordingly, it would be advantageous to provide a technique for network switching and data export responsive to message flow patterns.

SUMMARY OF THE INVENTION

The invention provides a method and system for switching in networks responsive to message flow patterns. A message "flow" is defined to comprise a set of packets to be transmitted between a particular source and a particular destination. When routers in a network identify a new

2

message flow, they determine the proper processing for packets in that message flow and cache that information for that message flow. Thereafter, when routers in a network identify a packet which is part of that message flow, they process that packet according to the proper processing for packets in that message flow. The proper processing may include a determination of a destination port for routing those packets and a determination of whether access control permits routing those packets to their indicated destination.

In another aspect of the invention, information about message flow patterns is collected, responsive to identified message flows and their packets. The collected information is reported to devices on the network. The collected information is used for a variety of purposes, including: to diagnose actual or potential network problems, to determine patterns of usage by date and time or by location, to determine which services and which users use a relatively larger or smaller amount of network resources, to determine which services are accessed by particular users, to determine which users access particular services, or to determine usage which falls within selected parameters (such as: access during particular dates or times, access to prohibited services, excessive access to particular services, excessive use of network resources, or lack of proper access).

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a network in which routing responsive to message flow patterns is performed.

FIG. 2 shows a method for routing in networks responsive to message flow patterns.

FIG. 3 shows data structures for use with a method for routing in networks responsive to message flow patterns.

FIG. 4 shows an IP address cache for use with a method for routing in networks responsive to message flow patterns.

FIG. 5 shows a method for collecting and reporting information about message flow patterns.

DESCRIPTION OF THE PREFERRED EMBODIMENT

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. However, those skilled in the art would recognize, after perusal of this application, that embodiments of the invention may be implemented using a set of general purpose computers operating under program control, and that modification of a set of general purpose computers to implement the process steps and data structures described herein would not require undue invention.

Message Flows

FIG. 1 shows a network in which routing responsive to message flow patterns is performed.

A network **100** includes at least one communication link **110**, at least one source device **120**, at least one destination device **130**, and at least one routing device **140**. The routing device **140** is disposed for receiving a set of packets **150** from the source device **120** and routing them to the destination device **130**.

The communication link **110** may comprise any form of physical media layer, such as ethernet, FDDI, or HDLC serial link.

The routing device **140** comprises a routing processor for performing the process steps described herein, and may include specific hardware constructed or programmed per-

3

forming the process steps described herein, a general purpose processor operating under program control, or some combination thereof.

A message flow **160** consists of a unidirectional stream of packets **150** to be transmitted between particular pairs of transport service access points (thus, network-layer addresses and port numbers). In a broad sense, a message flow **160** thus refers to a communication "circuit" between communication endpoints. In a preferred embodiment, a message flow **160** is defined by a network-layer address for a particular source device **120**, a particular port number at the source device **120**, a network-layer address for a particular destination device **130**, a particular port number at the destination device **130**, and a particular transmission protocol type. For example, the transmission protocol type may identify a known transmission protocol, such as UDP, TCP, ICMP, or IGMP (internet group management protocol).

In a preferred embodiment for use with a network of networks (an "internet"), the particular source device **120** is identified by its IP (internet protocol) address. The particular port number at the source device **120** is identified by either a port number which is specific to a particular process, or by a standard port number for the particular transmission protocol type. For example, a standard port number for the TCP protocol type is 6 and a standard port number for the UDP protocol type is 17. Other protocols which may have standard port numbers include the FTP protocol, the TELNET protocol, an internet telephone protocol, or an internet video protocol such as the "CUSEMe" protocol; these protocols are known in the art of networking. Similarly, the particular destination device **130** is identified by its IP (internet protocol) address; the particular port number at the destination device **130** is identified by either a port number which is specific to a particular process, or a standard port number for the particular transmission protocol type.

It will be clear to those skilled in the art, after perusing this application, that the concept of a message flow is quite broad, and encompasses a wide variety of possible alternatives within the scope and spirit of the invention. For example, in alternative embodiments, a message flow may be bi-directional instead of unidirectional, a message flow may be identified at a different protocol layer level than that of transport service access points, or a message flow may be identified responsive to other factors. These other factors may include one or more of the following: information in packet headers, packet length, time of packet transmission, or routing conditions on the network (such as relative network congestion or administrative policies with regard to routing and transmission).

Network Flow Switching

FIG. 2 shows a method for routing in networks responsive to message flow patterns.

In broad overview, the method for routing in networks responsive to message flow patterns comprises two parts. In a first part, the routing device **140** builds and uses a flow cache (described in further detail with regard to FIG. 3), in which routing information to be used for packets **150** in each particular message flow **160** is recorded and from which such routing information is retrieved for use. In a second part, the routing device **140** maintains the flow cache, such as by removing entries for message flows **160** which are no longer considered valid.

A method **200** for routing in networks responsive to message flow patterns is performed by the routing device **140**.

4

At a flow point **210**, the routing device **140** is disposed for building and using the flow cache.

At a step **221**, the routing device **140** receives a packet **150**.

At a step **222**, the routing device **140** identifies a message flow **160** for the packet **150**. In a preferred embodiment, the routing device **140** examines a header for the packet **150** and identifies the IP address for the source device **120**, the IP address for the destination device **130**, and the protocol type for the packet **150**. The routing device **140** determines the port number for the source device **120** and the port number for the destination device **130** responsive to the protocol type. Responsive to this set of information, the routing device **140** determines a flow key **310** (described with reference to FIG. 3) for the message flow **160**.

At a step **223**, the routing device **140** performs a lookup in a flow cache for the identified message flow **160**. If the lookup is unsuccessful, the identified message flow **160** is a "new" message flow **160**, and the routing device **140** continues with the step **224**. If the lookup is successful, the identified message flow **160** is an "old" message flow **160**, and the routing device **140** continues with the step **225**.

In a preferred embodiment, the routing device **140** determines a hash table key responsive to the flow key **310**. This aspect of the step **223** is described in further detail with regard to FIG. 3.

At a step **224**, the routing device **140** builds a new entry in the flow cache. The routing device **140** determines proper treatment of packets **150** in the message flow **160** and enters information regarding such proper treatment in a data structure pointed to by the new entry in the flow cache. In a preferred embodiment, the routing device **140** determines the proper treatment by performing a lookup in an IP address cache as shown in FIG. 4.

In a preferred embodiment, the proper treatment of packets **150** in the message flow **160** includes treatment with regard to switching (thus, the routing device **140** determines an output port for switching packets **150** in the message flow **160**), with regard to access control (thus, the routing device **140** determines whether packets **150** in the message flow **160** meet the requirements of access control, as defined by access control lists in force at the routing device **140**), with regard to accounting (thus, the routing device **140** creates an accounting record for the message flow **160**), with regard to encryption (thus, the routing device **140** determines encryption treatment for packets **150** in the message flow **160**), and any special treatment for packets **150** in the message flow **160**.

In a preferred embodiment, the routing device **140** performs any special processing for new message flows **160** at this time. For example, in one preferred embodiment, the routing device **140** requires that the source device **120** or the destination device **130** must authenticate the message flow **160**. In that case, the routing device **140** transmits one or more packets **150** to the source device **120** or the destination device **130** to request information (such as a user identifier and a password) to authenticate the new message flow **160**, and receives one or more packets **150** comprising the authentication information. This technique could be useful for implementing security "firewalls" and other authentication systems.

Thereafter, the routing device **140** proceeds with the step **225**, using the information from the new entry in the flow cache, just as if the identified message flow **160** were an "old" message flow **160** and the lookup in a flow cache had been successful.

5

At a step 225, the routing device 140 retrieves routing information from the entry in the flow cache for the identified message flow 160.

In a preferred embodiment, the entry in the flow cache includes a pointer to a rewrite function for at least part of a header for the packet 150. If this pointer is non-null, the routing device 140 invokes the rewrite function to alter the header for the packet 150.

At a step 226, the routing device 140 routes the packet 150 responsive to the routing information retrieved at the step 225.

Thus, in a preferred embodiment, the routing device 140 does not separately determine, for each packet 150 in the message flow 160, the information stored in the entry in the flow cache. Rather, when routing a packet 150 in the message flow 160, the routing device 140 reads the information from the entry in the flow cache and treats the packet 150 according to the information in the entry in the flow cache.

Thus, in a preferred embodiment, the routing device 140 routes the packet 150 to an output port, determines whether access is allowed for the packet 150, determines encryption treatment for the packet 150, and performs any special treatment for the packet 150, all responsive to information in the entry in the flow cache.

In a preferred embodiment, the routing device 140 also enters accounting information in the entry in the flow cache for the packet 150. When routing each packet 150 in the message flow 160, the routing device 140 records the cumulative number of packets 150 and the cumulative number of bytes for the message flow 160.

Because the routing device 140 processes each packet 150 in the message flow 160 responsive to the entry for the message flow 160 in the flow cache, the routing device 140 is able to implement administrative policies which are designated for each message flow 160 rather than for each packet 150. For example, the routing device 140 is able to reserve specific amounts of bandwidth for particular message flows 160 and to queue packets 150 for transmission responsive to the bandwidth reserved for their particular message flows 160.

Because the routing device 140 is able to associate each packet 150 with a particular message flow 160 and to associate each message flow 160 with particular network-layer source and destination addresses, the routing device 140 is able to associate network usage with particular workstations (and therefore with particular users) or with particular services available on the network. This can be used for accounting purposes, for enforcing administrative policies, or for providing usage information to interested parties.

For a first example, the routing device 140 is able to monitor and provide usage information regarding access using the HTTP protocol to world wide web pages at particular sites.

For a second example, the routing device 140 is able to monitor usage information regarding relative use of network resources, and to give priority to those message flows 160 which use relatively fewer network resources. This can occur when a first message flow 160 is using a relatively low-bandwidth transmission channel (such as a 28.8 kilobits per second modem transmission channel) and when a second message flow 160 is using a relatively high-bandwidth transmission channel (such as a T-1 transmission line).

At a flow point 230, the routing device 140 is disposed for maintaining the flow cache.

6

At a step 241, the routing device 140 examines each entry in the flow cache and compares a current time with a last time a packet 150 was routed using that particular entry. If the difference exceeds a first selected timeout, the message flow 160 represented by that entry is considered to have expired due to nonuse and thus to no longer be valid.

In a preferred embodiment, the routing device 140 also examines the entry in the flow cache and compares a current time with a first time a packet 150 was routed using that particular entry. If the difference exceeds a second selected timeout, the message flow 160 represented by that entry is considered to have expired due to age and thus to no longer be valid. The second selected timeout is preferably about one minute.

Expiring message flows 160 due to age artificially requires that a new message flow 160 must be created for the next packet 150 in the same communication session represented by the old message flow 160 which was expired. However, it is considered preferable to do so because it allows information to be collected and reported about message flows 160 without having to wait for those message flows 160 to expire from nonuse. For example, a multiple-broadcast communication session could reasonably last well beyond the time message flows 160 are expired for age, and if not so expired would mean that information about network usage would not account for significant network usage.

In a preferred embodiment, the routing device 140 also examines the entry in the flow cache and determines if the "next hop" information has changed. If so, the message flow 160 is expired due to changed conditions. Other changed conditions which might cause a message flow 160 to be expired include changes in access control lists or other changes which might affect the proper treatment of packets 150 in the message flow 160. The routing device 140 also expires entries in the flow cache on a least-recently-used basis if the flow cache becomes too full.

If the message flow 160 is still valid, the routing device 140 continues with the next entry in the flow cache until all entries have been examined. If the message flow 160 is no longer valid, the routing device 140 continues with the step 242.

At a step 242, the routing device 140 collects historical information about the message flow 160 from the entry in the flow cache, and deletes the entry.

Flow Cache

FIG. 3 shows data structures for use with a method for routing in networks responsive to message flow patterns.

A flow cache 300 comprises a memory which associates flow keys 310 with information about message flows 160 identified by those flow keys 310. The flow cache 300 includes a set of buckets 301. Each bucket 301 includes a linked list of entries 302. Each entry 302 includes information about a particular message flow 160, including routing, access control, accounting, special treatment for packets 150 in that particular message flow 160, and a pointer to information about treatment of packets 150 to the destination device 130 for that message flow 160.

In a preferred embodiment, the flow cache 300 includes a relatively large number of buckets 301 (preferably about 16,384 buckets 301), so as to minimize the number of entries 302 per bucket 301 and thus so as to minimize the number of memory accesses per entry 302. Each bucket 301 comprises a four-byte pointer to a linked list of entries 302. The linked list preferably includes only about one or two entries 302 at the most.

7

In a preferred embodiment, each entry **302** includes a set of routing information, a set of access control information, a set of special treatment information, and a set of accounting information, for packets **150** in the message flow **160**.

The routing information comprises the output port for routing packets **150** in the message flow **160**.

The access control information comprises whether access is permitted for packets **150** in the message flow **160**.

The accounting information comprises a time stamp for the first packet **150** in the message flow **160**, a time stamp for the most recent packet **150** in the message flow **160**, a cumulative count for the number of packets **150** in the message flow **160**, and a cumulative count for the number of bytes **150** in the message flow **160**.

IP Address Cache

FIG. 4 shows an IP address cache for use with a method for routing in networks responsive to message flow patterns.

An IP address cache **400** comprises a tree having a root node **410**, a plurality of inferior nodes **410**, and a plurality of leaf data structures **420**.

Each node **410** comprises a node/leaf indicator **411** and an array **412** of pointers **413**.

The node/leaf indicator **411** indicates whether the node **410** is a node **410** or a leaf data structure **420**; for nodes **410** it is set to a "node" value, while for leaf data structures **420** it is set to a "leaf" value.

The array **412** has room for exactly 256 pointers **413**; thus, the IP address cache **400** comprises an M-trie with a branching width of 256 at each level. M-tries are known in the art of tree structures. IP addresses comprise four bytes, each having eight bits and therefore 256 possible values. Thus, each possible IP address can be stored in the IP address cache **400** using at most four pointers **413**.

The inventors have discovered that IP addresses in actual use are unexpectedly clustered, so that the size of the IP address cache **400** is substantially less, by a factor of about five to a factor of about ten, than would be expected for a set of randomly generated four-byte IP addresses.

Each pointer **413** represents a subtree of the IP address cache **400** for its particular location in the array **412**. Thus, for the root node **410**, the pointer **413** at location **3** represents IP addresses having the form 3.xxx.xxx.xxx, where "xxx" represents any possible value from zero to 255. Similarly, in a subtree for IP addresses having the form 3.xxx.xxx.xxx, the pointer **413** at location **141** represents IP addresses having the form 3.141.xxx.xxx. Similarly, in a subtree for IP addresses having the form 3.141.xxx.xxx, the pointer **413** at location **59** represents IP addresses having the form 3.141.59.xxx. Similarly, in a subtree for IP addresses having the form 3.141.59.xxx, the pointer **413** at location **26** represents the IP address 3.141.59.26.

Each pointer **413** is either null, to indicate that there are no IP addresses for the indicated subtree, or points to an inferior node **410** or leaf data structure **420**. A least significant bit of each pointer **413** is reserved to indicate the type of the pointed-to structure; that is, whether the pointed-to structure is a node **410** or a leaf data structure **420**. In a preferred embodiment where pointers **413** must identify an address which is aligned on a four-byte boundary, the two least significant bits of each pointer **413** are unused for addressing, and reserving the least significant bit for this purpose does not reduce the scope of the pointer **413**.

Each leaf data structure comprises information about the IP address, stored in the IP address cache **400**. In a preferred embodiment this information includes the proper processing for packets **150** addressed to that IP address, such as a determination of a destination port for routing those packets

8

and a determination of whether access control permits routing those packets to their indicated destination.

Flow Data Export

FIG. 5 shows a method for collecting and reporting information about message flow patterns.

A method **500** for collecting and reporting information about message flow patterns is performed by the routing device **140**.

At a flow point **510**, the routing device **140** is disposed for obtaining information about a message flow **160**. For example, in a preferred embodiment, as noted herein, the routing device **140** obtains historical information about a message flow **160** in the step **242**. In alternative embodiments, the routing device **140** may obtain information about message flows **160**, either in addition or instead, by occasional review of entries in the flow cache, or by directly monitoring packets **150** in message flows **160**.

It will be clear to those skilled in the art, after perusing this application, that the concept of reporting information about message flows is quite broad, and encompasses a wide variety of possible alternatives within the scope and spirit of the invention. For example, in alternative embodiments, information about message flows may include bi-directional traffic information instead of unidirectional traffic information, information about message flows may include information at a different protocol layer level other than that of transport service access points and other than that at which the message flow is itself defined, or information about message flows may include actual data transmitted as part of the message flow itself. These actual data may include one or more of the following: information in packet headers, information about files of file names transmitted during the message flow, or usage conditions of the message flow (such as whether the message flow involves steady or bursty transmission of data, or is relatively interactive or relatively unidirectional).

At a step **521**, the routing device **140** obtains historical information about a particular message flow **160**, and records that information in a flow data table.

At a step **522**, the routing device **140** determines a size of the flow data table, and compares that size with a selected size value. If the flow data table exceeds the selected size value, the routing device **140** continues with the step **523** to report flow data. If the flow data table does not exceed the selected size value, the routing device **140** returns to the step **521** to obtain historical information about a next particular message flow **160**.

At a step **523**, the routing device **140** builds an information packet, responsive to the information about message flows **160** which is recorded in the flow data table.

At a step **524**, the routing device **140** transmits the information packet to a selected destination device **130** on the network **100**. In a preferred embodiment, the selected destination device **130** is determined by an operating parameter of the routing device **140**. This operating parameter is set when the routing device **140** is initially configured, and may be altered by an operator of the routing device **140**.

In a preferred embodiment, the selected destination device **130** receives the information packet and builds (or updates) a database in the format for the RMON protocol. The RMON protocol is known in the art of network monitoring.

At a flow point **530**, a reporting device **540** on the network **100** is disposed for reporting using information about message flows **160**.

At a step **531**, the reporting device **540** queries the selected destination device **130** for information about mes-

sage flows **160**. In a preferred embodiment, the reporting device **540** uses the RMON protocol to query the selected destination device **130** and to obtain information about message flows **160**.

At a step **532**, the reporting device **540** builds a report about a condition of the network **100**, responsive to information about message flows **160**.

At a step **533**, the reporting device **540** displays or transmits that report about the condition of the network **100** to interested parties.

In preferred embodiments, the report may comprise one or more of a wide variety of information, and interested parties may use that information for one or more of a wide variety of purposes. Some possible purposes are noted herein:

Interested parties may diagnose actual or potential network problems. For example, the report may comprise information about packets **150** in particular message flows **160**, including a time stamp for a first packet **150** and a time stamp for a last packet **150** in the message flow **160**, a cumulative total number of bytes in the message flow **160**, a cumulative total number of packets **150** in the message flow **160**, or other information relevant to diagnosing actual or potential network problems.

Interested parties may determine patterns of usage of the network by date and time or by location. For example, the report may comprise information about which users or which services on the network are making relatively heavy use of resources. In a preferred embodiment, usage of the network **100** is displayed in a graphical form which shows use of the network **100** in a false-color map, so that network administrators and other interested parties may rapidly determine which services, which users, and which communication links are relatively loaded or relatively unloaded with demand.

Interested parties may determine which services are accessed by particular users, or which users access particular services. For example, the report may comprise information about which services are accessed by particular users at a particular device on the network **100**, or which users access a particular service at a particular device on the network **100**. This information may be used to market or otherwise enhance these services. In a preferred embodiment, users who access a particular world wide web page using the HTTP protocol are recorded, and information is sent to those users about changes to that web page and about further services available from the producers of that web page. Providers of the particular web page may also collect information about access to their web page in response to date and time of access, and location of accessing user.

Information about patterns of usage of the network, or about which services are accessed by particular users, or which users access particular services, may be used to implement accounting or billing for resources, or to set limits for resource usage, such as by particular users, by particular service providers, or by particular protocol types (and therefore by particular types of services).

Interested parties may determine usage which falls within (or without) selected parameters. These selected parameters may involve access during particular dates or times, such as for example access to particular services during or outside normal working hours. For example, it may be desirable to record those accesses to a company database which occur outside normal working hours.

These selected parameters may involve access to prohibited services, excessive access to particular services, or excessive use of network resources, such as for example access to particular servers using the HTTP protocol or the FTP protocol which fall within (or without) a particular administrative policy. For example, it may be desirable to

record accesses to repositories of games or other recreational material, particularly those accesses which occur within normal working hours.

These selected parameters may involve or lack of proper access, such as for example access control list failures or unauthorized attempts to access secure services. For example, it may be desirable to record unauthorized attempts to access secure services, particularly those attempts which form a pattern which might indicate a concerted attempt to gain unauthorized access.

In alternative embodiments, the routing device **140** could save the actual packets **150** for the message flow **160**, or some part thereof, for later examination. For example, a TELNET session (a message flow **160** comprising use of the TELNET protocol by a user and a host) could be recorded in its entirety, or some portion thereof, for later examination, e.g., to diagnose problems noted with the network or with the particular host.

In further alternative embodiments, the routing device **140** could save the actual packets **150** for selected message flows **160** which meet certain selected parameters, such as repeated unauthorized attempts to gain access.

In embodiments where actual packets **150** of the message flow **160** are saved, it would be desirable to perform a name translation (such as a reverse DNS lookup), because the IP addresses for the source device **120** and the destination device **130** are transitory. Thus, it would be preferable to determine the symbolic names for the source device **120** and the destination device **130** from the IP addresses, so that the recorded data would have greater meaning at a later time.

Alternative Embodiments

Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to those skilled in the art after perusal of this application.

We claim:

1. A method in a router for determining how to treat packets, comprising the steps of:

determining a flow and a flow key for a packet;

performing a first lookup in a flow cache based on the flow key, with the flow determined to be old if the first lookup is successful and determined to be new if the first lookup is unsuccessful;

if the flow is old, using results of the first lookup to determine how to treat the packet; and

if the flow is new, performing a second lookup in an IP address cache, with results of the second lookup used to determine how to treat the packet and to make an entry for the flow in the flow cache such that if a next packet in the flow is received, the flow is identified as old, and the next packet is treated in a same manner as the packet.

2. A method as in claim 1, wherein the flow key is determined based on a source IP address and port, destination IP address and port, and protocol type for the packet.

3. A method as in claim 2, wherein determining how to treat the packet includes determining how to route the packet from the router.

4. A method as in claim 3, wherein determining how to treat the packet further includes determining whether access is allowed for the packet.

5. A method as in claim 3, wherein determining how to treat the packet further includes determining encryption treatment for the packet.

6. A method as in claim 1, further comprising the step of entering accounting information in the flow cache for the flow.

11

7. A method as in claim 1, wherein the IP address cache is an M-trie having a branching width of at least 16, said M-trie having a set of nodes including a root node, each said node having a pointer and a node/leaf indicator included within a data word, said node/leaf indicator determining whether said pointer points to an inferior node or to a leaf data structure, each leaf data structure including information associated with one IP address, said pointer to said inferior node being located at an offset in an array of nodes, said offset within an Nth level of said M-trie corresponding to an associated value represented by an Nth byte of said IP address, wherein because said set of IP addresses are clustered, said M-trie is smaller, by a factor of about five to a factor of about ten, than if said set of IP addresses were randomly distributed.

8. A router that determines how to treat packets, comprising:

links for receiving packets from at least one source device and routing those packets to at least one destination device;

a processor for executing instructions; and

a memory for storing data and said instruction, said data including a flow cache and an IP address cache, and said instructions including the steps of (a) determining a flow and a flow key for a packet, (b) performing a first lookup in the flow cache based on the flow key, with the flow determined to be old if the first lookup is successful and determined to be new if the first lookup is unsuccessful, (c) if the flow is old, using results of the first lookup to determine how to treat the packet, and (d) if the flow is new, performing a second lookup in the IP address cache, with results of the second lookup used to determine how to treat the packet and to make an entry for the flow in the flow cache such that if a next packet in the flow is received, the flow is identified as old and the next packet is treated in the same manner as the packet.

9. A router as in claim 8, wherein the flow key is determined based on a source IP address and port, destination IP address and port, and protocol type for the packet.

10. A router as in claim 9, wherein determining how to treat the packet includes determining how to route the packet from the router.

11. A router as in claim 10, wherein determining how to treat the packet further includes determining whether access is allowed for the packet.

12. A router as in claim 10, wherein determining how to treat the packet further includes determining encryption treatment for the packet.

13. A router as in claim 8, wherein the instructions further comprise the step of entering accounting information in the flow cache for the flow.

14. A router as in claim 8, wherein the IP address cache is an M-trie having a branching width of at least 16, said M-trie having a set of nodes including a root node, each said node having a pointer and a node/leaf indicator included within a data word, said node/leaf indicator determining whether said pointer points to an inferior node or to a leaf data structure, each leaf data structure including information associated with one IP address, said pointer to said inferior node being located at an offset in an array of nodes, said offset within an Nth level of said M-trie corresponding to an associated value represented by an Nth byte of said IP address, wherein because said set of IP addresses are clustered, said M-trie is smaller, by a factor of about five to a factor of about ten, than if said set of IP addresses were randomly distributed.

12

15. A memory storing instructions for causing a router to determine how to treat packets, the instructions comprising the steps of:

determining a flow and a flow key for a packet;

performing a first lookup in a flow cache based on the flow key, with the flow determined to be old if the first lookup is successful and determined to be new if the first lookup is unsuccessful;

if the flow is old, using results of the first lookup to determine how to treat the packet; and

if the flow is new, performing a second lookup in an IP address cache, with results of the second lookup used to determine how to treat the packet and to make an entry for the flow in the flow cache such that if a next packet in the flow is received, the flow is identified as old and the next packet is treated in the same manner as the packet.

16. A memory as in claim 15, wherein the flow key is determined based on a source IP address and port, destination IP address and port, and protocol type for the packet.

17. A memory as in claim 16, wherein determining how to treat the packet includes determining how to route the packet from the router.

18. A memory as in claim 17, wherein determining how to treat the packet further includes determining whether access is allowed for the packet.

19. A memory as in claim 17, wherein determining how to treat the packet further includes determining encryption treatment for the packet.

20. A memory as in claim 15, wherein the instructions further comprise the step of entering accounting information in the flow cache for the flow.

21. A memory as in claim 15, wherein the IP address cache is an M-trie having a branching width of at least 16, said M-trie having a set of nodes including a root node, each said node having a pointer and a node/leaf indicator included within a data word, said node/leaf indicator determining whether said pointer points to an inferior node or to a leaf data structure, each leaf data structure including information associated with one IP address, said pointer to said inferior node being located at an offset in an array of nodes, said offset within an Nth level of said M-trie corresponding to an associated value represented by an Nth byte of said IP address, wherein because said set of IP addresses are clustered, said M-trie is smaller, by a factor of about five to a factor of about ten, than if said set of IP addresses were randomly distributed.

22. A router that determines how to treat packets, comprising:

means for determining a flow and a flow key for a packet; means for performing a first lookup in a flow cache based on the flow key, with the flow determined to be old if the first lookup is successful and determined to be new if the first lookup is unsuccessful;

means for using, if the flow is old, results of the first lookup to determine how to treat the packet; and

means for performing, if the flow is new, a second lookup in an IP address cache, with results of the second lookup used to determine how to treat the packet and to make an entry for the flow in the cache such that if a next packet in the flow is received, the flow is identified as old, and the next packet is treated in a same manner as the packet.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,889,181 B2
DATED : May 3, 2005
INVENTOR(S) : Darren R. Kerr et al.

Page 1 of 2

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

Title page.

Item [56], **References Cited**, U.S. PATENT DOCUMENTS, insert the following:

-- 4,933,938	6/1990	Sheehy
5,006,982	4/1991	Ebersole et al.
5,267,235	11/1993	Thacker
5,287,535	2/1994	Sakagawa et al.
5,291,442	3/1994	Emma et al.
5,325,504	6/1994	Tipley et al.
5,347,642	9/1994	Barratt
5,394,408	2/1995	Nishihara et al.
5,396,491	3/1995	Newman
5,418,922	5/1995	Liu
5,426,637	6/1995	Derby et al.
5,444,491	8/1995	Lim
5,450,406	9/1995	Esaki et al.
5,515,376	5/1996	Murthy et al.
5,523,999	6/1996	Takano et al.
5,524,254	6/1996	Morgan et al.
5,528,592	6/1996	Schibler et al.
5,533,033	7/1996	Ratner
5,557,747	9/1996	Rogers et al.
5,583,865	12/1996	Esaki et al.
5,608,908	3/1997	Barghouti et al.
5,610,904	3/1997	Eng et al.
5,614,891	3/1997	Zeinstra et al.
5,621,720	4/1997	Bronte et al.
5,644,751	7/1997	Burnett
5,754,768	5/1998	Brech et al.
5,774,660 *	6/1998	Brendel et al.
5,842,040	11/1998	Hughes et al.
5,862,136	1/1999	Irwin
5,892,924 *	4/1999	Lyon et al.
5,925,097	7/1999	Gopinath et al.
6,091,725 *	7/2000	Cheriton et al.
6,182,139 B1 *	1/2001	Brendel
6,219,699 B1 *	4/2001	Mc Cloghrie et al.
6,308,148 B1 *	10/2001	Bruins et al.
6,343,322 B2	1/2002	Nagami et al. --.

FOREIGN PATENT DOCUMENTS, insert the following:

-- WO	96/34479 A1	10/1996
WO	99/17182 A2	4/1999 --.

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,889,181 B2
DATED : May 3, 2005
INVENTOR(S) : Darren R. Kerr et al.

Page 2 of 2

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

Title page (cont'd).

OTHER PUBLICATIONS, insert the following:

-- Sangoma Technologies, Inc., "Frame Relay Packet Switching," Frame Relay Tutorial, 1999, 3 pages.

Nicholas William McKeown, et al., "Scheduling Algorithms For Input Quented Cells Switches," 1995, 118 pages.

Bruce A. Mah, "Measurements And Observations Of IP Multicast Traffic," pages 1-12.

Nick McKeown, et al., "Achieving 100% Throughout In An Input Queed Switch" pages 1-22.

Nick McKeown, et al., "The Baybridge: A High Speed Bridge/Router," Presented at IFIP PFHSN Workshop, Stockholm, Sweden, May 1992, Project Report: 12, Revision: 2.0, pages 1-18.

Nick McKeown, et al., "Bridges, Routers, And Switches," 112 pages.

Nick McKeown, et al., "A Quantitative Comparison of Scheduling Algorithms For Input-Queed Switches" pages 1-16.

Pankaj Gupta, et al., "Scheduling Algorithms For An Input Queed Variable Length Packet Switch," Stanford University, pages 1-34.

Brent T. Jensen, "Access Networking", pages 1-8.

Cisco Systems, Inc., "Net Flow Services And Applications (White Paper)", 1999, pages 1-27. --.


Column 12,

Lines 38-41, delete "or to a leaf data structure, each leaf data structure including information associated with one IP address, said pointer to said inferior node".

Line 43, delete "Nth byte" and insert -- Nth Byte --.

Signed and Sealed this

Thirty-first Day of January, 2006



JON W. DUDAS

Director of the United States Patent and Trademark Office