US005974503A

United States Patent [19]

Venkatesh et al.

[11] **Patent Number: 5,974,503**

[45] **Date of Patent: Oct. 26, 1999**

[54] **STORAGE AND ACCESS OF CONTINUOUS MEDIA FILES INDEXED AS LISTS OF RAID STRIPE SETS ASSOCIATED WITH FILE NAMES**

[75] Inventors: **Dinesh Venkatesh**, Brookline; **John Forecast**, Newton; **Wayne W. Duso**, Shrewsbury, all of Mass.

[73] Assignee: **EMC Corporation**, Hopkinton, Mass.

[21] Appl. No.: **08/851,509**

[22] Filed: **May 5, 1997**

Related U.S. Application Data

[60] Provisional application No. 60/044,948, Apr. 25, 1997.

[51] **Int. Cl.**[6] **G06F 12/00**

[52] **U.S. Cl.** **711/114**; 395/182.04; 395/182.05

[58] **Field of Search** 711/114; 371/51.1, 371/40.4; 395/182.04, 182.05

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,608,688	8/1986	Hansen et al.	371/11
4,755,928	7/1988	Johnson et al.	364/200
4,761,785	8/1988	Clark et al.	371/51.1
5,072,378	12/1991	Manka	395/575
5,148,432	9/1992	Gordon et al.	371/10.1
5,175,837	12/1992	Arnold et al.	395/425
5,206,939	4/1993	Yani et al.	395/400
5,208,665	5/1993	McCalley et al.	358/862
5,208,813	5/1993	Stallmo et al.	371/10.1
5,218,695	6/1993	Noveck et al.	395/600

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

0 061 570 A3	10/1982	European Pat. Off.	G11C 9/06
633694 A1	1/1995	European Pat. Off.	H04N 7/173
0 683 464 A2	11/1995	European Pat. Off.	G06F 17/60
0 697 660 A1	2/1996	European Pat. Off.	G06F 11/20
WO 93/16557	8/1993	WIPO	H04N 7/14
WO 94/00816	1/1994	WIPO	G06F 13/00
WO 95/10918	4/1995	WIPO	H04N 7/14
WO 97/16023	5/1997	WIPO	H04N 7/173

OTHER PUBLICATIONS

Dishon et al., "A Highly Available Storage System Using the Checksum Method," Seventeenth International Symposium on Fault–Tolerant Computing, Jul. 6–8, 1987, Pittsburg, Penn., IEEE Computer Society, IEEE, New York, N.Y., pp. 178–181.

Martin E. Schulze, "Considerations in the Design of a RAID Prototype," Computer Science Division, EECS, University of California at Berkeley, CA, Aug. 25, 1988.

Katz et al., "Disk System Architectures for High Performance Computing," Computer Science Division, EECS, University of California at Berkeley, CA, Mar. 1989.

(List continued on next page.)

Primary Examiner—John W. Cabeca
Assistant Examiner—Nasser Moazzami
Attorney, Agent, or Firm—Arnold White & Durkee

[57] **ABSTRACT**

A continuous media file is comprised of stripe sets over disk drives in one or more RAID sets. In a preferred embodiment, the RAID set includes n disk drives. The data storage of each disk drive in the RAID set is partitioned into an integer number m of hyper-volumes, and the parity is stored in one hyper-volume of each of m disk drives in the RAID set. The stripe set includes a series of transfer units of data in respective ones of the disk drives. Each transfer unit includes an integer number j of data blocks, and each hyper-volume includes an integer number k of transfer units. Each stripe set includes (m)(n−1) transfer units of data. The transfer units of the RAID set are allocated for the storage of continuous media data in a right-to-left and then top-to-bottom order in which the transfer units appear in an m row by n column matrix in which the rows of the matrix represent parity groups of hyper-volumes in the disk drives and the columns of the matrix represent storage in the respective disk drives. At most one write access to each parity hyper-volume need be performed during write access to a stripe set. Parity changes for the data being written are accumulated in non-volatile memory, and written to the RAID set after completion of the writing of the data.

12 Claims, 36 Drawing Sheets

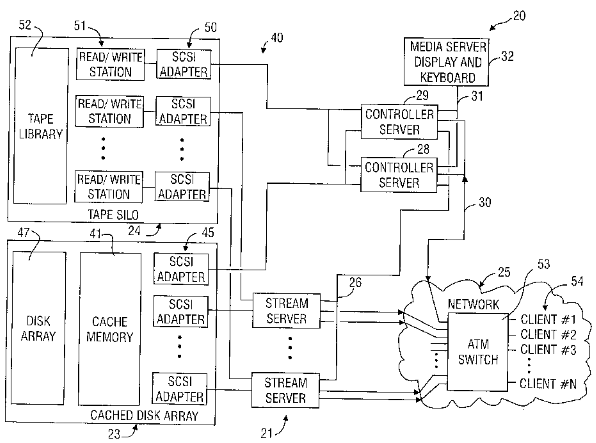

U.S. PATENT DOCUMENTS

5,235,601	8/1993	Stallmo et al.	371/40.4
5,255,270	10/1993	Yanai et al.	371/10.2
5,263,145	11/1993	Brady et al.	395/425
5,269,011	12/1993	Yanai et al.	395/425
5,274,799	12/1993	Brant et al.	385/575
5,276,860	1/1994	Fortier et al.	395/575
5,276,867	1/1994	Kenley et al.	395/600
5,335,352	8/1994	Yanai et al.	395/800
5,341,493	8/1994	Yanai et al.	395/425
5,367,698	11/1994	Webber et al.	395/800
5,371,532	12/1994	Gelman et al.	348/7
5,381,539	1/1995	Yanai et al.	395/425
5,408,465	4/1995	Gusella et al.	370/17
5,410,343	4/1995	Coddington et al.	348/7
5,414,455	5/1995	Hooper et al.	348/7
5,442,390	8/1995	Hooper et al.	348/7
5,442,749	8/1995	Northcutt et al.	395/200.09
5,442,771	8/1995	Filepp et al.	395/200.08
5,477,263	12/1995	O'Callaghan et al.	348/7
5,508,733	4/1996	Kassatly	348/13
5,519,435	5/1996	Anderson	348/8
5,528,282	6/1996	Voeten et al.	348/7
5,528,513	6/1996	Vaitzblit et al.	364/514 A
5,530,557	6/1996	Asit et al.	358/342
5,533,021	7/1996	Branstad et al.	370/60.1
5,534,912	7/1996	Kostreski	348/6
5,537,408	7/1996	Branstad et al.	370/79
5,537,534	7/1996	Voigt et al.	395/182.04
5,539,660	7/1996	Blair et al.	364/514 C
5,544,313	8/1996	Shachnai et al.	395/200.01
5,544,327	8/1996	Dan et al.	395/250
5,544,345	8/1996	Carpenter et al.	395/477
5,544,347	8/1996	Yanai et al.	395/489
5,550,577	8/1996	Verbiest et al.	348/7
5,550,982	8/1996	Long et al.	395/200.13
5,553,005	9/1996	Voeten et al.	364/514 R
5,557,317	9/1996	Nishio et al.	348/7
5,559,949	9/1996	Reimer et al.	395/161
5,572,660	11/1996	Jones	395/182.04
5,574,662	11/1996	Windrem et al.	364/514 R
5,579,475	11/1996	Blaum et al.	395/182.05
5,583,561	12/1996	Baker et al.	347/7
5,586,264	12/1996	Belknap et al.	395/200.08
5,594,910	1/1997	Filepp et al.	395/800
5,603,058	2/1997	Belknap et al.	395/855
5,606,359	2/1997	Youden et al.	348/7
5,625,405	4/1997	DuLac et al.	348/7
5,633,810	5/1997	Mandal et al.	364/514
5,657,439	8/1997	Jones et al.	395/182.05
5,812,753	9/1998	Chiariotti	395/182.04

OTHER PUBLICATIONS

Gray et al., "Parity Striping of Disc Arrays: Low–Cost Reliable Storage with Acceptable Throughput," Technical Report 90.2, Tandem Computers, Cupertino, CA., Jan. 1990.

Edward K. Lee, "Software and Performance Issues in the Implementation of a RAID Prototype," Computer Science Division, EECS, University of California at Berkeley, CA, Mar. 1989.

Sincoskie, WD, "System architecture for a large scale video on demand service," Computer Networks and ISDN Systems, vol. 22, No. 2, Sep. 1991, pp. 155–162.

Tobagi FA, Pang J, "StarWorks (Trademark)—A video applications server," Proceedings, IEEE COMPCON 93, San Francisco, Calif., 1993, pp. 4–11.

Vaitzblit L, "The design and implementation of a high bandwidth file service for continuous media," Master's Thesis, Massachusetts Institute of Technology, Cambridge, Mass., Nov. 4, 1991.

Patterson et al., "A Case for Redundant Arrays Of Inexpensive Disks (RAID)," Report No. UCB/CSD 87/391, Computer Science Division (EECS), University of California, Berkeley, California, Dec. 1987, pp. 1–24.

Patterson et al., "Introduction to Redundant Arrays of Inexpensive Disks (RAID)," COMPCON 89 Proceedings, Feb. 27–Mar. 3, 1989, IEEE Computer Society, pp. 112–117.

Ousterhout et al., "Beating the I/O Bottleneck: A Case for Log–Structured File Systems," Operating Systems Review, vol. 23, No. 1, ACM Press, Jan., 1989, pp. 11–28.

Douglis et al., "Log Structured File Systems," COMPCON 89 Proceedings, Feb. 27–Mar. 3, 1989, IEEE Computer Society, pp. 124–129.

Rosemblum et al., "The Design and Implementation of a Log–Structured File System," ACM Transactions on Computer Systems, vol. 1, Feb. 1992, pp. 26–52.

M(aurice) William Collins, "A Network File Storage System," IEEE Seventh Symposium on Mass Storage Systems, Nov. 4–7, 1985, Tuscon, AZ, pp. 1–11, Los Alamos Nat. Lab. No. LA–UR–85–3183.

John H. Howard, "An Overview of the Andrew File System," USENIX Winter Conference, Feb. 9–12, 1988, Dallas, TX, p. 23–26.

John H. Howard et al., "Scale and Performance in a Distributed File System," ACM Transactions on Computer Systems, vol. 6, No. 1, Feb. 1988, p. 51–81.

David C. Steere et al., "Efficient User–Level File Cache Management on the Sun Vnode Interface," USENIX Summer Conference, Jun. 11–15, 1990, Anaheim, CA, p. 325–331.

Matt Blaze et al., "Long–Term Caching Strategies for Very Large Distributed File Systems," USENIX, Summer '91, Nashville, TN, p. 3–15.

Thomas W. Page, Jr., et al., "Management of Replicated Volume Location Data in the Ficus Replicated File System," USENIX, Summer '91, Nashville, TN, p. 17–29.

Storage Computer Corporation, "High Performance, Fault Tolerant Disk Array Platform For File Servers And Computer Systems," 1991, Nashua, NH 12 pages.

France Telecom, "Telesauvegarde," 26 page paper dated Nov. 2, 1994 about work based on France Telecom patent application "Dispositif et Proceed de Sauvegaude a Distance de Donnees Numeriques" [Method and Apparatus for Safeguarding at a Distance Numeric Data], Institute National de la Propiete Industrielle, France, Appln. No. 93.12771, Oct. 28, 1994 (filed Oct. 21, 1993), and partial English translation (15 pages).

Krishnan Natarajan, "Video Servers Take Root," IEEE Spectrum, Apr. 1995, IEEE, New York, NY, pp. 66–69.

K. K. Ramakrishnan et al., "Operating system support for a video–on–demand file service," Multimedia Systems (1995) 3:53–65.

Ralf Steinmetz, "Analyzing the Multimedia Operating System," IEEE MultiMedia, Spring 1995, pp. 68–84.

Audrey Chou, "EMC, Computer–Storage Leader, Still Hears Footsteps," The Wall Street Journal, Aug. 9, 1995, Dow Jones & Co., Princeton, N.J.

Pardhu Vadlamudi, "EMC Hunts for Solution to Bottlenecks," InfoWorld, Apr. 15, 1996, #1590, San Mateo, CA 94402.

Michael Goldberg, "EMC to Pump Data Over Networks," *Computerworld,* Apr. 15, 1996.

"EMC Moves Into Servers," *Broadcasting Cable,* Apr. 15, 1996.

"Symmetrix Model 55XX Product Manual, P/N 200–810–550 Rev D," EMC Corporation, Hopkinton, Mass., May 1994, pp. 1–236.

"NFS: Network File System Protocol Specification," RFC 1094, Sun Microsystems, Inc., Mar. 1989, pp. 1–27.

J. Case, M. Fedor, M. Schoffstall, J. Davin, "A Simple Network Management Protocol (SNMP)," May 1990, MIT Laboratory for Computer Science, Cambridge Mass., pp. 1–35.

Rangen PV, Vin HM, "Designing file systems for digital video and audio," Proceedings of the 13th ACM Symposium on Operating Systems Principles, Monterey, Calif., 1991, pp. 81–94.

Vin HM, Rangan PV, (1993) "Designing a multiuser HDTV storage server," IEEE Journal on Selected Areas in Communications, vol. 11, No. 1, Jan. 1993, pp. 153–164.

Anderson DP, Osawa Y, Govindan R, "A file system for continuous media," ACM Transactions on Computer Systems, vol. 10., No. 4, Nov. 1992, pp. 311–337.

Federighi C, "A Distributed Hierarchical Storage Manager for a Video–on–Demand System," Department of Electrical Engr. and Computer Science, University of California, Berkeley, California, Dec. 1993.

Haskins R, "The shark continuous–media file server," Proceedings IEEE COMPCOM 93, San Francisco, Calif., 1993, pp. 12–15.

Little TD, Rhanger G, Folz RJ, Gibbon JF, Reeve FW, Schelleng DH, Venkatesh D, "A digital on–demand video service supporting content based queries," Proceedings of ACM Multimedia 93, Anaheim, Calif., Aug. 1–6, 1993, pp. 427–436.

Lougher, P. Sheperd, D. "The design of a storage server for continuous media," The Computer Journal, vol. 36, No. 1, 1993, pp. 32–42.

Rangan PV, Vin HM, Ramanathan S, "Designing an on–demand multimedia service," IEEE Communications Magazine, vol. 30, No. 7, Jul. 1992, pp. 56–64.

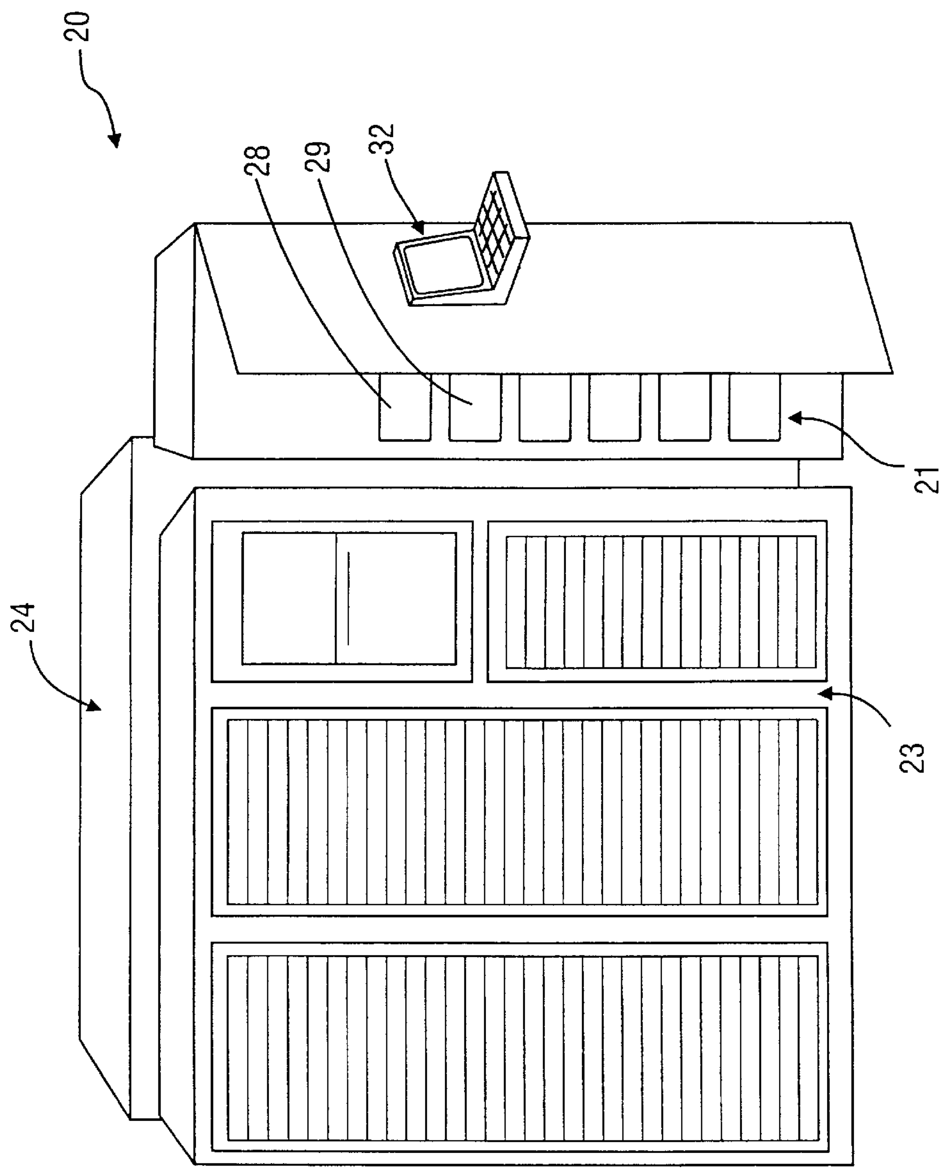

FIG. 1

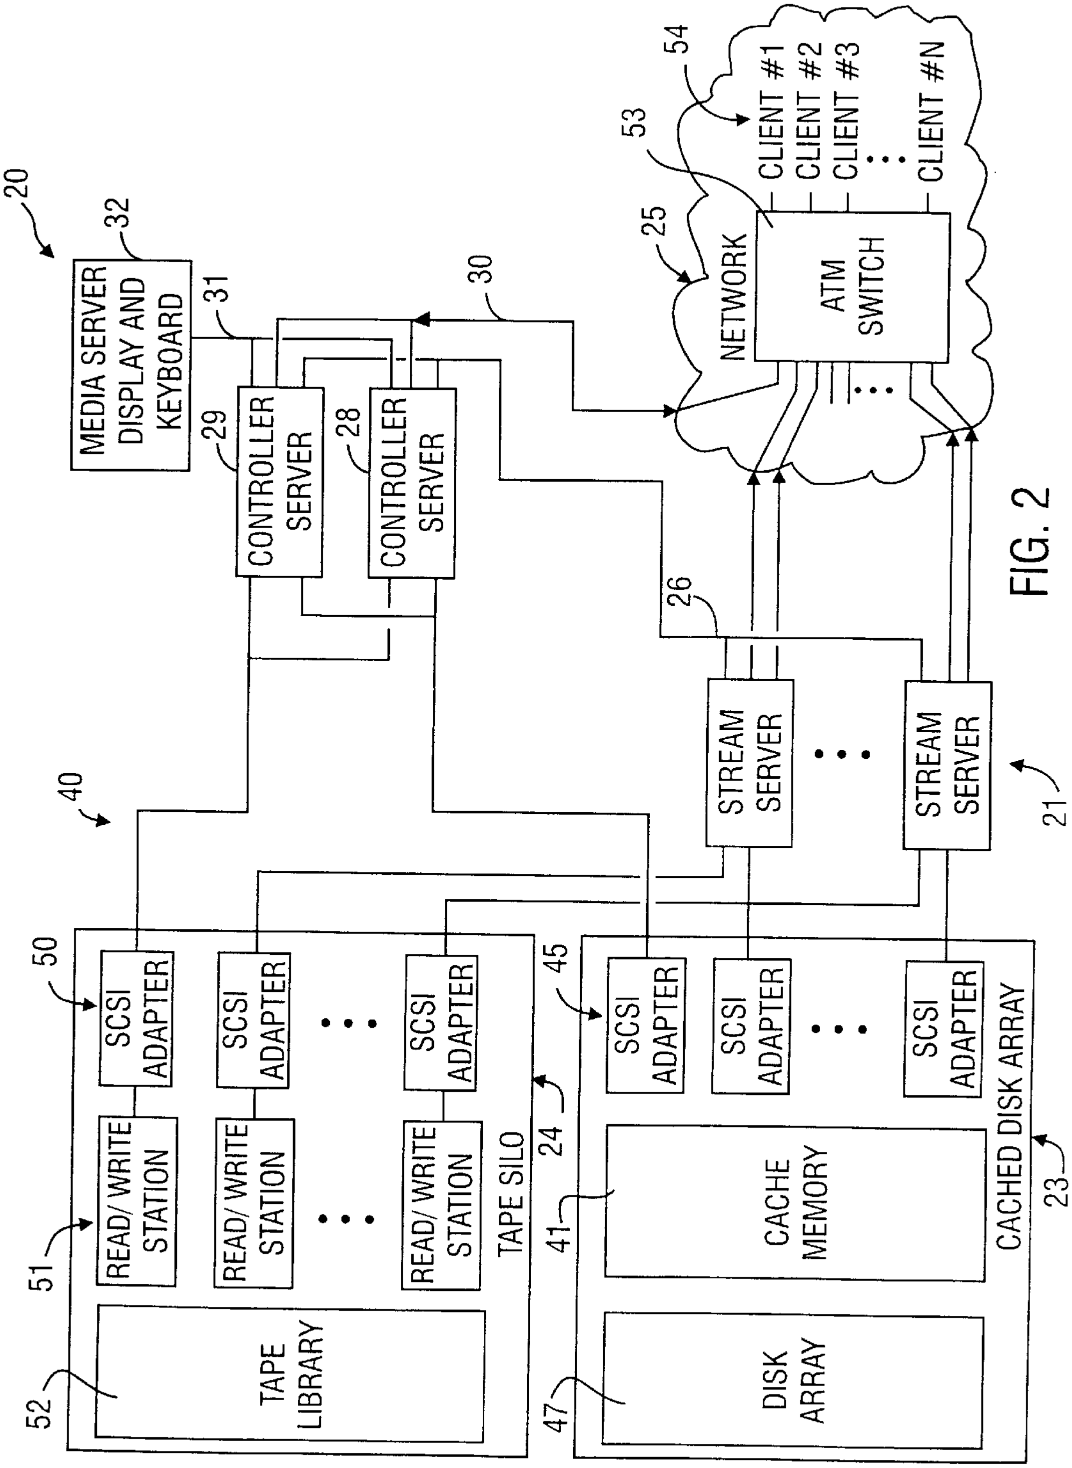
20
MEDIA SERVER DISPLAY AND KEYBOARD
32
31
29
CONTROLLER SERVER
28
CONTROLLER SERVER
30
25
NETWORK
ATM SWITCH
53
54
CLIENT #1
CLIENT #2
CLIENT #3
CLIENT #N
26
STREAM SERVER
STREAM SERVER
21
40
50
SCSI ADAPTER
SCSI ADAPTER
SCSI ADAPTER
51
READ/WRITE STATION
READ/WRITE STATION
READ/WRITE STATION
TAPE SILO
24
52
TAPE LIBRARY
45
SCSI ADAPTER
SCSI ADAPTER
SCSI ADAPTER
41
CACHE MEMORY
47
DISK ARRAY
CACHED DISK ARRAY
23

FIG. 2

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | sales@docketalarm.com | 1-866-77-FASTCASE

fastcase®
Smarter legal research.