UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
12/553,199	09/03/2009	Kamleshkumar K. Lad	060692-8065.US01	6112

116504 7590 10/28/2016

Perkins Coie LLP - CommVault Systems, Inc.
P.O. Box 1247
Seattle, WA 98111-1247

EXAMINER
HUTTON, NAN

ART UNIT	PAPER NUMBER
2154	

NOTIFICATION DATE	DELIVERY MODE
10/28/2016	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

patentprocurement@perkinscoie.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte KAMLESHKUMAR K. LAD[1]

Appeal 2015-003099
Application 12/553,199
Technology Center 2100

Before JAMES W. DEJMEK, SCOTT B. HOWARD, and MATTHEW J. McNEILL, *Administrative Patent Judges*.

DEJMEK, *Administrative Patent Judge*.

DECISION ON APPEAL

Appellant appeals under 35 U.S.C. § 134(a) from a Final Rejection of claims 1, 2, 4, 5, 7–12, and 14–18. Claims 3, 6, 13, and 19–22 have been canceled. App. Br. 3. We have jurisdiction over the remaining pending claims under 35 U.S.C. § 6(b).

We affirm.

---

[1] Appellant identifies CommVault Systems, Inc. as the real party in interest. App. Br. 2.

STATEMENT OF THE CASE

Introduction

Appellant's claimed invention is directed to transferring or migrating data objects (such as files, folders, data stores, and/or discrete data components) "by migrating segments, portions, increments, or proper subsets of the data objects." Spec. ¶ 16. Exemplary approaches described include block-based and chunk-based migration methodologies. Spec. ¶¶ 36–39, 48–50. In disclosed embodiments, only a proper subset (i.e., less than all) of blocks or chunks is migrated, based on predetermined storage criteria (e.g., last time the data block was modified). Spec. ¶¶ 17–18, 36–39, 48–50, Figs. 6A, 6B. According to the Specification, the claimed invention provides cost and time benefits, *inter alia*, associated with data storage and restoration. Spec. ¶¶ 39, 64.

Claim 1 is representative of the subject matter on appeal and is reproduced below with the disputed limitations emphasized in *italics*:

> 1. A system, comprising a storage device, for migrating data from a primary storage device to a secondary storage device, wherein the system includes a file system for transferring data to the primary storage device, and wherein the system further includes a disk driver for at least writing data received from the file system to the primary storage device and a secondary driver for at least writing data to the secondary storage device, the system comprising:
>
> a virtual disk driver that is between the file system and the primary storage device, is capable of mounting to the file system, and receives data from the file system associated with the primary storage device and provides data to the disk driver that writes data to the primary storage device, wherein the virtual disk driver includes:
>
> a data reception component, wherein the data reception component is configured to receive data from the file system, *wherein the received data identifies multiple blocks of a file to be*

modified, wherein the multiple blocks are a proper subset of the total number of blocks for the file;

a data interception component, wherein the data interception component is configured to intercept the received data and extract information associated with the received data, wherein the extracted information includes information identifying the multiple blocks to be modified;

an index component, wherein the index component is configured to update an index that associates the extracted information with data blocks on the secondary storage device that contain the received data; and

a data transfer component, wherein the data transfer component is configured to transfer the received data to the secondary driver for storage to the secondary storage device;

a block-level data migration component, wherein the block-level migration component is configured to identify data blocks within the primary storage device that satisfy one or more predetermined criteria;

a data management component, configured to communicate with the virtual disk driver, the block-level data migration component and one or more media agents, wherein the data management component includes a storage policy that provides the one or more predetermined criteria, the storage policy identifying a time period in which to retain data within the primary storage device and identifying the one or more media agents in which to transfer the data from the file system to the disk driver, via the virtual disk driver;

a media agent, wherein the media agent is one of the identified media agents and is configured to:

transfer data from the identified data blocks to the secondary driver, while retaining remaining data blocks on the primary storage device; and

update an index that associates the transferred data with the secondary storage device that stores data from the secondary driver; and

a chunking component configured to divide data already in the secondary storage device into chunks based on logical divisions

including a date of creation, deletion or reception or a size of data to be chunked;

wherein the block-level data migration component is configured to identify data blocks that have not changed after a predetermined time period based on data from the file system, and

wherein the multiple blocks are written to the secondary storage device in a format that is not native to a format for an application that created the file.

The Examiner's Rejections[2]

1. Claims 1 and 2 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Chakravarty et al. (US 2007/0208788 A1; Sept. 6, 2007) ("Chakravarty") and Minami et al. (US 2007/0214330 A1; Sept. 13, 2007) ("Minami"). Final Act. 3–8.

2. Claims 4, 5, 7–12, and 14–18 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Chakravarty, Minami, and Reisman (US 7,751,628 B1; July 6, 2010). Final Act. 8–19.

Issues on Appeal

1. Did the Examiner err in finding Chakravarty teaches or suggests identifying or transferring less than all of the data blocks comprising a data object to a secondary storage, as recited in independent claim 1 ("a proper subset of the total number of blocks for the file"), and as commensurately recited in independent claim 5 ("some data blocks being

[2] In response to Appellant's amendment of claim 11, the Examiner has withdrawn the rejection of claims 11, 12, and 14–18 under 35 U.S.C. § 101 as being directed to non-statutory subject matter. Adv. Act. 2 (mailed March 11, 2014).

retained on the first data store while the identified blocks being transferred to the second data store") and independent claim 11 ("transferring . . . no more than n-1 data blocks of the data file from the primary storage to the secondary storage")?

2. Did the Examiner err in relying on Chakravarty in rejecting claim 5 because the Examiner incorrectly equates Chakravarty's "data blocks" with the claimed "data objects" recited in claim 5?

3. Did the Examiner fail to provide articulated reasoning with rational underpinning to support the proposed combination of Chakravarty and Minami with Reisman?

ANALYSIS[3]

Claims 1 and 2

Although Appellant concedes Chakravarty teaches subdividing datasets into data blocks (App. Br. 9), Appellant contends Chakravarty fails to teach or fairly suggest that only a proper subset (i.e., less than all) of the data blocks for a particular dataset are identified and transferred to a secondary storage while the remaining data blocks are retained in primary storage. App. Br. 8–12; Reply Br. 4–6. In particular, Appellant argues "Chakravarty teaches that all of the dataset is transferred from a primary storage to a secondary storage." App. Br. 9.

[3] Throughout this Decision, we have considered the Appeal Brief, filed July 22, 2014 ("App. Br."); the Reply Brief, filed December 23, 2014 ("Reply Br."); the Examiner's Answer, mailed on October 23, 2014 ("Ans."); and the Final Office Action ("Final Act."), mailed on November 22, 2013, from which this Appeal is taken.

In response, the Examiner finds Chakravarty teaches that one or more data blocks are moved between the first tier (i.e., primary) and second tier (i.e., secondary) storage devices based on a data migration parameter according to a user selected access policy. Ans. 3 (citing Chakravarty ¶¶ 33, 34, Fig. 1); *see also* Final Act. 4, 6 (citing Chakravarty ¶¶ 4, 31, 33–35).

Figure 1 of Chakravarty is illustrative and is reproduced below:

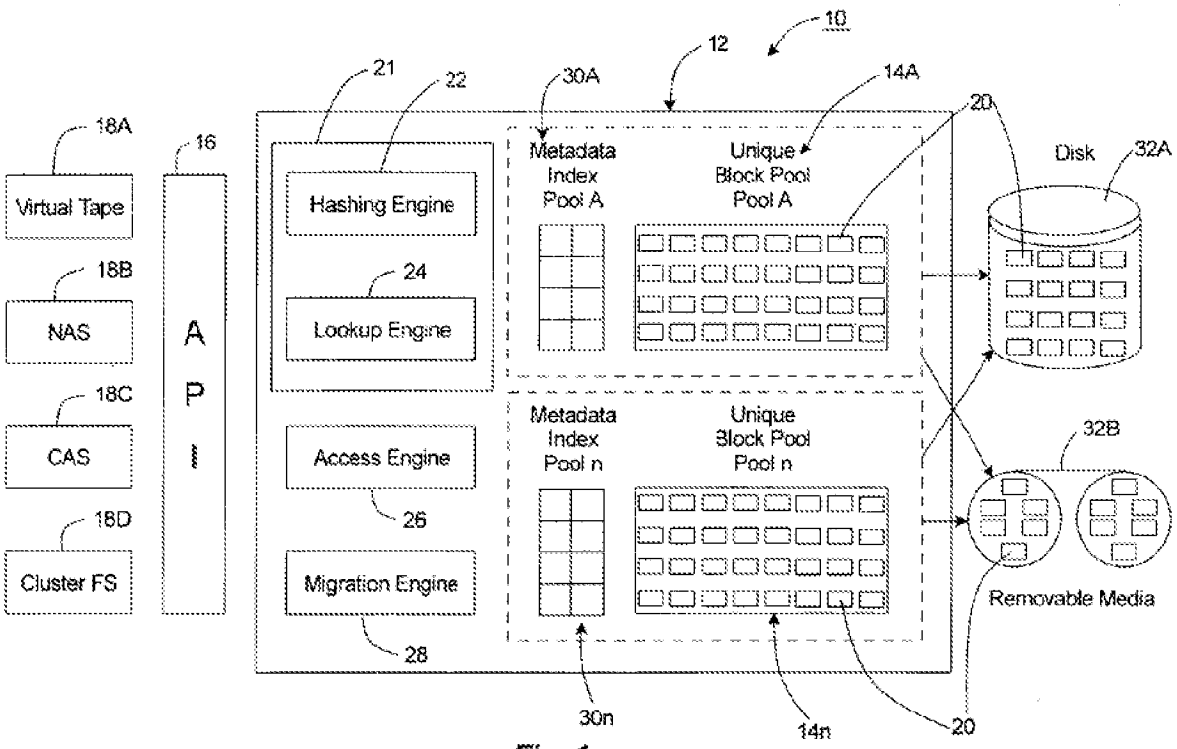

Fig. 1

Figure 1 of Chakravarty is a block diagram of the claimed data storage system. Chakravarty ¶ 9. As shown, Chakravarty's data storage system (10) comprises unique data block pools (14A–14n) and a redundancy reducer (21). Chakravarty ¶ 24. "The redundancy reducer **21** utilizes the hash engine **22** and the lookup engine **24** *to store only unique data blocks **20*** in the unique data block pools **14**A–**14**n. By storing only the unique data

blocks **20**, a substantial savings of storage space can be realized."

Chakravarty ¶ 24 (emphasis added). Additionally, Chakravarty teaches:

> In some embodiments, the first tier storage device includes a disk drive *that stores at least a portion of the data set* prior to reduction of the redundancy of the data blocks by the redundancy reducer. Further, the second tier storage device can include a disk drive that *stores at least a portion of the data set* following reduction of the redundancy of the data blocks by the redundancy reducer.

Chakravarty ¶ 6 (emphases added). Thus, Chakravarty teaches, *inter alia*, identifying a proper subset of the total number of data blocks of a data set and storing or transferring a portion of the data blocks (i.e., a proper subset) to secondary storage.

In the Reply Brief, Appellant contends the Examiner's reliance on paragraph 34 of Chakravarty undercuts the Examiner's position. Reply Br. 4–5. In particular, Appellant argues Chakravarty teaches calculating a hash handle on the entire data set, and not on a proper subset. Reply Br. 4–5 (citing Chakravarty ¶¶ 33, 34).

We do not find Appellant's argument persuasive of Examiner error. Rather, Chakravarty teaches the creation of a pointer map wherein the various pointers point to corresponding data blocks. *See* Chakravarty ¶ 33.

For the reasons discussed *supra*, we are unpersuaded of Examiner error. Accordingly, we sustain the Examiner's rejection of independent claim 1 and of claim 2, which depends therefrom and was not argued separately. App. Br. 13, 18.

Claims 4, 5 and 7–10

Independent claim 5 is similar to claim 1, but is directed to a data object (as opposed to a file) and recites in relevant part "some data blocks being retained on the first data store while the identified data blocks being transferred to the second data store." The blocks to be transferred are identified by a bitmap which identifies data blocks of a data object that satisfy a recent access storage criterion. App. Br. 13.

Appellant contends Chakravarty tracks access to de-duplicated data blocks but not when individual data sets have been accessed. App. Br. 13; Reply Br. 6–8. Appellant advances a similar argument as was argued with respect to claim 1 that Chakravarty does not teach transferring a portion of the identified data blocks to a second data store. App. Br. 14.

For similar reasons as discussed with respect to claim 1, we find this argument unpersuasive of Examiner error. Further, the Examiner finds, and we agree, Chakravarty teaches a manager residing between the file systems and storage system (i.e., an intermediate component) that maintains a pointer map that tracks information about changes or updates to the data blocks of the data sets. Final Act. 10 (citing Chakravarty ¶¶ 21–35). Chakravarty expressly teaches that one of the migration criteria may be the last time the block was accessed. Chakravarty ¶ 35.

Appellant additionally argues the Examiner improperly equates Chakravarty's data blocks with the claimed data objects. App. Br. 14–15; Reply Br. 6–8. Contrary to Appellant's "important distinction" of Chakravarty tracking when data blocks have been accessed as opposed to when the data set has been accessed, we are unpersuaded of error. Chakravarty teaches that the data set is comprised of data blocks. *See*

Chakravarty ¶ 4. Thus, when an individual data block has been accessed, a person of ordinary skill in the art would understand that the data set to which the data block belongs has also been accessed. *See* Ans. 6–7.

Additionally, Appellant argues the Examiner "failed to provide adequate rationale why one of skill in the art would have looked to Reisman to modify the teaching of Chakravarty." App. Br. 18–19; Reply Br. 8–10. In particular, Appellant asserts Reisman is directed to deleting data, whereas Chakravarty (as well as Appellant's claimed invention) is directed to preserving data. App. Br. 18.

The Federal Circuit has stated that "rejections on obviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness." *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006), *cited with approval in KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 418 (2007).

As the Examiner explains, Reisman is concerned with managing storage space. Ans. 7. As part of the disclosed system, Reisman teaches a file allocation table to track used and free storage blocks to allow the system to delete low-significance data without opening the whole file. Ans. 7 (citing Reisman, col. 8, ll. 1–20). The Examiner further explains:

> Both systems [(Chakravarty and Reisman)] need to analyze data in files using the allocation table, Chakravarty migrates low-significance data to a lower tier storage device; while Reisman deletes low-significance data. The difference of the two references is the intended use of the technique, and is not relevant in determining the reason for combining the two references.

Ans. 8.

Thus, we find the Examiner has provided an articulated reasoning to support the proposed combination of references and conclusion of obviousness. *See* KSR, 550 U.S. at 418. Accordingly, we are unpersuaded the Examiner erred in relying on the combination of Chakravarty, Minami, and Reisman.

For the reasons discussed *supra*, we are unpersuaded of Examiner error. Accordingly, we sustain the Examiner's rejection of independent claim 5 and of dependent claims 4 and 7–10, which were not argued separately. App. Br. 18.

Claims 11, 12, and 14–18

Independent Claim 11 recites, *inter alia,* "transferring data contained by the identified no more than n-1 data blocks of the data file from the primary storage to the secondary storage, but retaining remaining blocks of the data file in the primary storage." Similar to claim 1, Appellant argues Chakravarty fails to teach identifying n-1 data blocks that represent a portion of a data file consisting of n blocks (i.e., identifying less than all, or a proper subset) and transferring the n-1 blocks to secondary storage. App. Br. 16–17.

For similar reasons as those discussed with respect to claim 1, we do not find Appellant's argument persuasive of Examiner error. Accordingly, we sustain the Examiner's rejection of independent claim 11 and of claims 12 and 14–18, which depend therefrom and were not argued separately. App. Br. 18.

DECISION

We affirm the Examiner's decision to reject claims 1, 2, 4, 5, 7–12, and 14–18.

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a). *See* 37 C.F.R. § 41.50(f).

AFFIRMED