

US010459917B2

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

(10) **Patent No.:** **US 10,459,917 B2**
(45) **Date of Patent:** ***Oct. 29, 2019**

(54) **PLUGGABLE STORAGE SYSTEM FOR DISTRIBUTED FILE SYSTEMS**

(58) **Field of Classification Search**
CPC G06F 17/30463; G06F 17/30224; G06F 17/30433; G06F 17/30445;
(Continued)

(71) Applicant: **EMC IP Holding Company LLC**,
Hopkinton, MA (US)

(72) Inventors: **Lei Chang**, Beijing (CN); **Tao Ma**,
Beijing (CN); **Zhanwei Wang**, Beijing
(CN); **Lirong Jian**, Beijing (CN); **Lili
Ma**, Beijing (CN); **Gavin Sherry**, San
Mateo, CA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,191,611 A 3/1993 Lang
5,495,607 A 2/1996 Pisello
(Continued)

(73) Assignee: **EMC IP Holding Company LLC**,
Hopkinton, MA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 680 days.

Liao et al. Multi-dimensional Index on Hadoop Distributed File
System, 2010, Fifth IEEE International Conference on Networking,
Architecture, and Storage, pp. 240-249.

This patent is subject to a terminal dis-
claimer.

Primary Examiner — Robert W Beausoliel, Jr.

Assistant Examiner — Susan F Rayyan

(21) Appl. No.: **15/150,263**

(74) *Attorney, Agent, or Firm* — Van Pelt, Yi & James
LLP

(22) Filed: **May 9, 2016**

(65) **Prior Publication Data**

US 2016/0292181 A1 Oct. 6, 2016

(57) **ABSTRACT**

Related U.S. Application Data

(63) Continuation of application No. 13/843,067, filed on
Mar. 15, 2013, now Pat. No. 9,454,548.

A method, article of manufacture, and apparatus for man-
aging data. In some embodiments, this includes an initial
instruction for a file stored in a first storage system, deter-
mining that the initial instruction is not supported by the first
storage system, identifying a combination of instructions to
the first storage system after determining that the initial
instruction is not supported by the first storage system,
wherein the combination of instructions is based on the
initial instruction, performing the identified combination of
instructions on the file stored in the first storage system, and
storing results of the performed identified combination of
instructions.

(Continued)

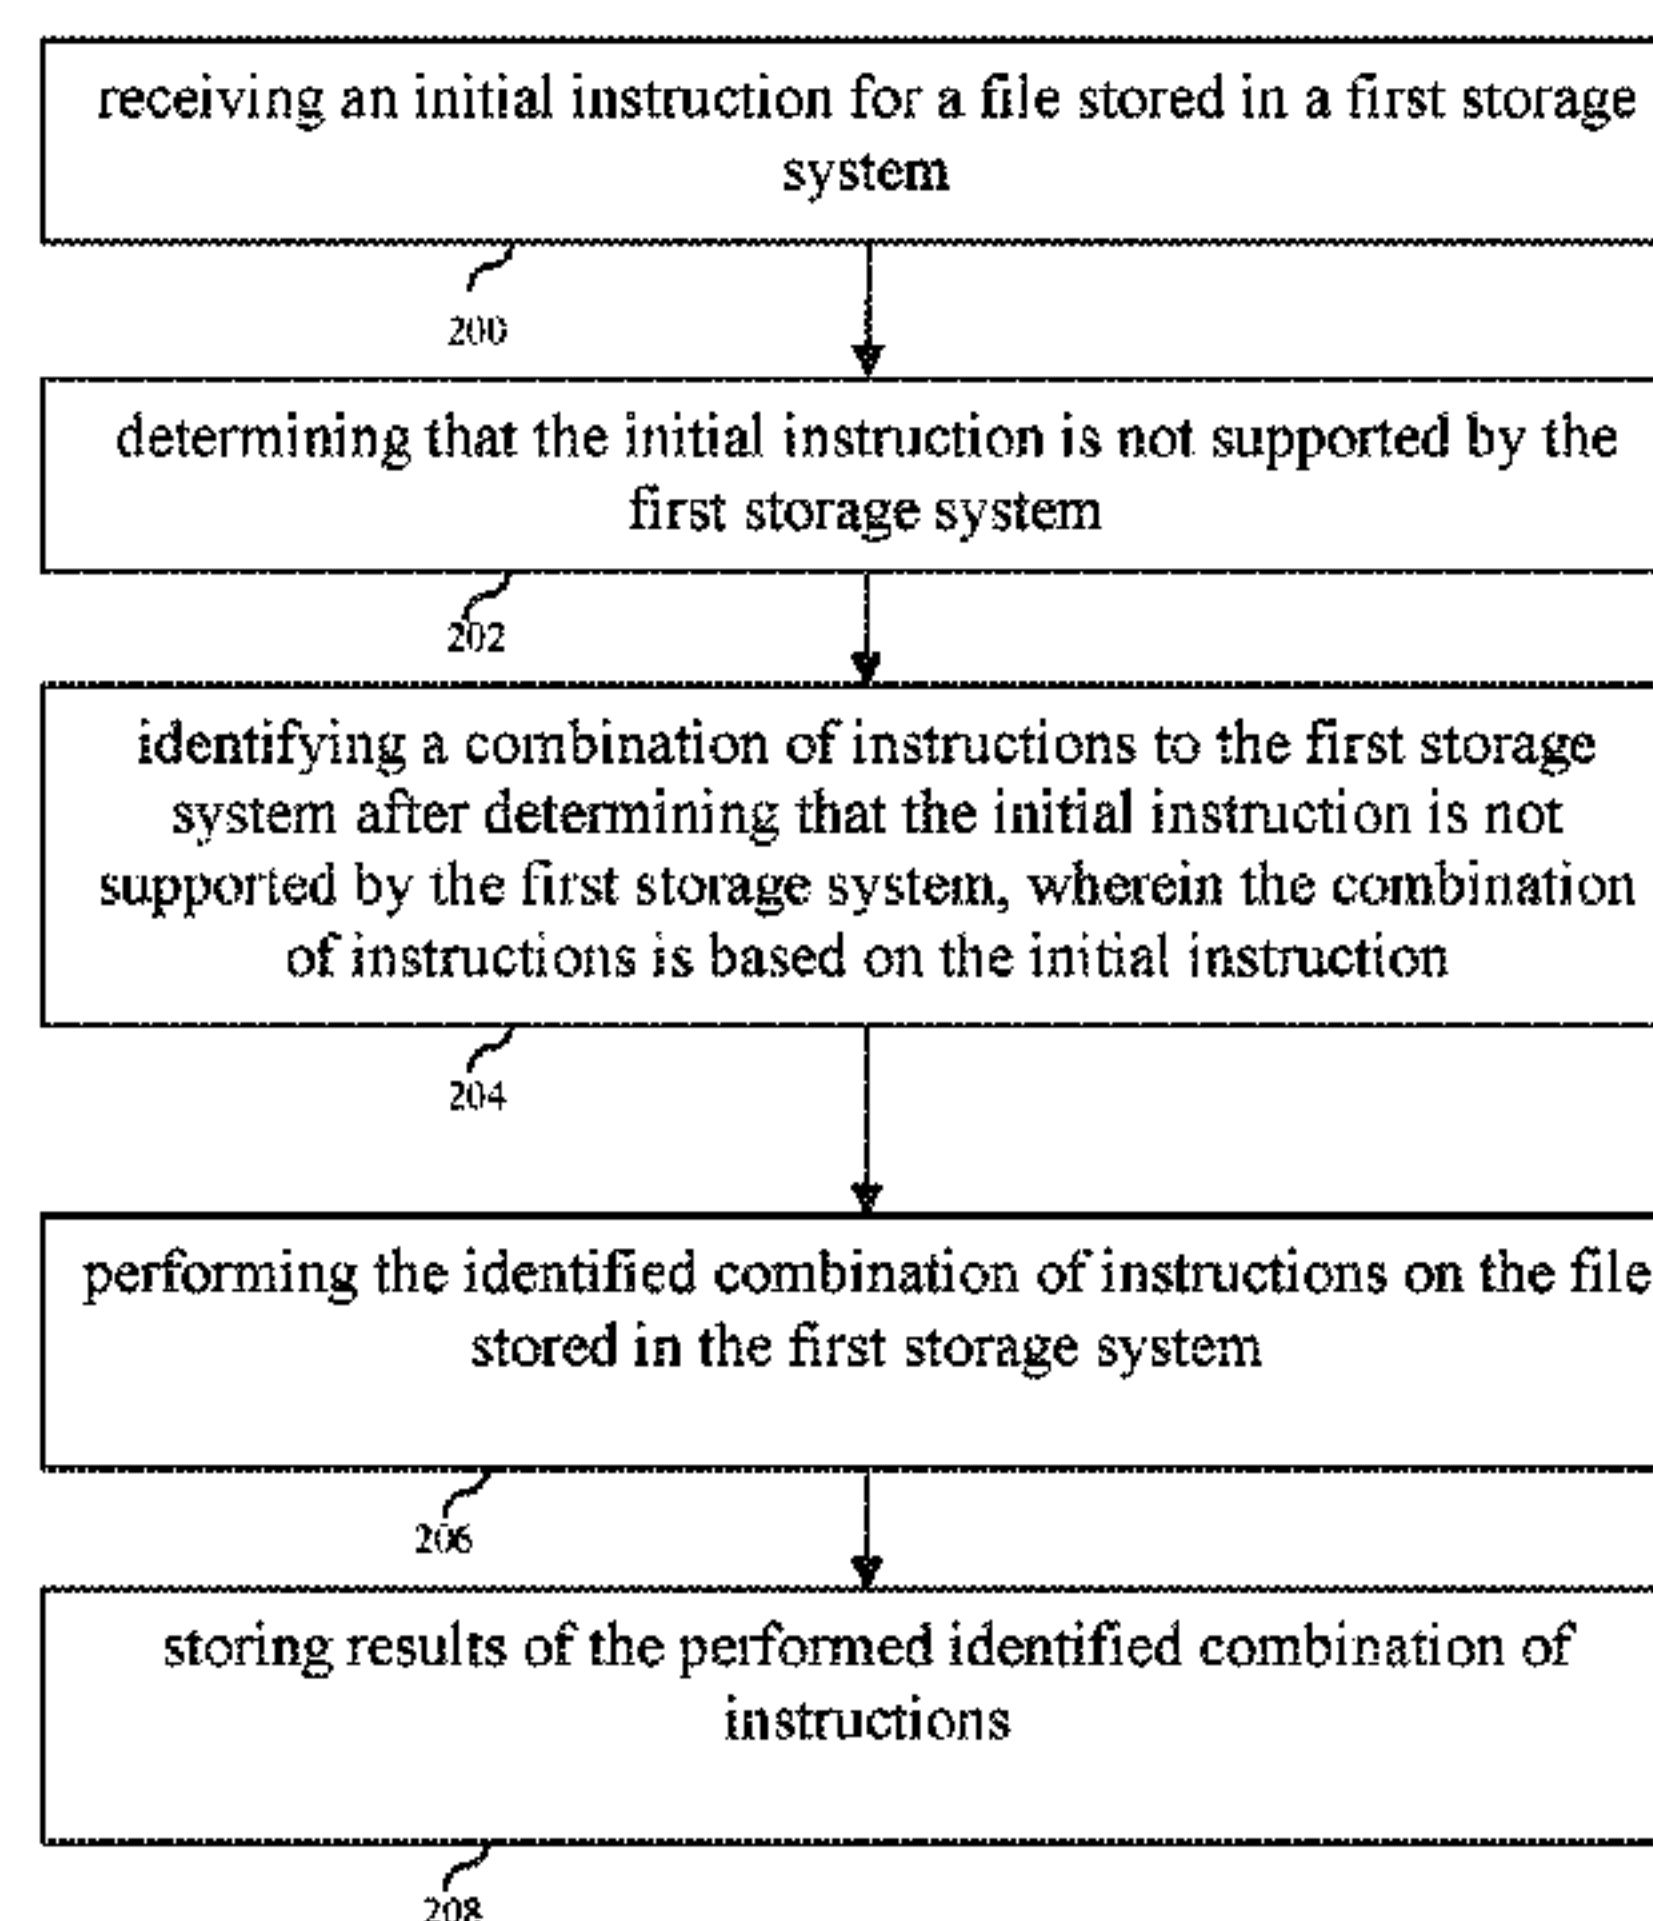
(51) **Int. Cl.**
G06F 16/2453 (2019.01)
G06F 16/10 (2019.01)

(Continued)

(52) **U.S. Cl.**
CPC **G06F 16/24542** (2019.01); **G06F 16/10**
(2019.01); **G06F 16/11** (2019.01);

(Continued)

20 Claims, 2 Drawing Sheets



Related U.S. Application Data						
		7,475,199	B1 *	1/2009	Bobbitt	G06F 16/119 711/154
(60)	Provisional application No. 61/769,043, filed on Feb. 25, 2013.	7,493,311	B1	2/2009	Cutsinger	
		7,593,938	B2	9/2009	Lemar	
		7,689,609	B2	3/2010	Lango	
		7,761,678	B1	7/2010	Bodmer	
(51)	Int. Cl.	7,774,335	B1	8/2010	Scofield	
	<i>G06F 16/11</i> (2019.01)	7,958,303	B2	6/2011	Shuster	
	<i>G06F 16/27</i> (2019.01)	7,984,043	B1	7/2011	Waas	
	<i>G06F 16/14</i> (2019.01)	8,010,738	B1	8/2011	Chilton	
	<i>G06F 16/182</i> (2019.01)	8,185,488	B2 *	5/2012	Chakravarty	G06N 5/022 706/47
	<i>G06F 16/907</i> (2019.01)	8,195,769	B2	6/2012	Miloushev et al.	
	<i>G06F 16/18</i> (2019.01)	8,332,526	B2 *	12/2012	Kruse	H04L 29/08072 709/230
	<i>G06F 16/2455</i> (2019.01)	8,417,681	B1	4/2013	Miloushev et al.	
	<i>G06F 16/2458</i> (2019.01)	8,484,259	B1	7/2013	Makkar et al.	
	<i>G06F 16/2452</i> (2019.01)	8,543,564	B2 *	9/2013	Conrad	G06F 17/30522 707/708
	<i>H04L 29/08</i> (2006.01)	8,578,096	B2	11/2013	Malige et al.	
	<i>H04L 29/06</i> (2006.01)	8,595,237	B1	11/2013	Chaudhary et al.	
	<i>G06F 16/43</i> (2019.01)	8,751,533	B1 *	6/2014	Dhavale	G06F 3/0647 707/782
	<i>G06F 16/21</i> (2019.01)	8,762,330	B1	6/2014	Kick	
	<i>G06F 16/245</i> (2019.01)	8,825,752	B1	9/2014	Madhavan	
(52)	U.S. Cl.	8,832,154	B1	9/2014	Srinivasan et al.	
	CPC	8,856,286	B2	10/2014	Barsness	
	<i>G06F 16/148</i> (2019.01); <i>G06F 16/182</i> (2019.01); <i>G06F 16/1858</i> (2019.01); <i>G06F 16/2453</i> (2019.01); <i>G06F 16/2455</i> (2019.01); <i>G06F 16/2471</i> (2019.01); <i>G06F 16/24524</i> (2019.01); <i>G06F 16/24532</i> (2019.01); <i>G06F 16/27</i> (2019.01); <i>G06F 16/907</i> (2019.01); <i>H04L 65/60</i> (2013.01); <i>H04L 67/1097</i> (2013.01); <i>H05K 999/99</i> (2013.01); <i>G06F 16/113</i> (2019.01); <i>G06F 16/217</i> (2019.01); <i>G06F 16/245</i> (2019.01); <i>G06F 16/43</i> (2019.01)	8,971,916	B1	3/2015	Joyce et al.	
		9,118,697	B1	8/2015	Kishore	
		9,146,766	B2 *	9/2015	Shaikh	G06F 12/0261
		9,323,758	B1	4/2016	Stacey et al.	
		9,449,007	B1	9/2016	Wood	
		9,674,294	B1 *	6/2017	Gonthier	H04L 69/08
		9,727,588	B1	8/2017	Ostapovicz	
		2002/0002638	A1	1/2002	Obara	
		2002/0049782	A1	4/2002	Herzenberg et al.	
		2002/0156840	A1 *	10/2002	Ulrich	G06F 16/10 709/203
		2003/0126120	A1	7/2003	Faybishenko	
(58)	Field of Classification Search	2003/0229637	A1	12/2003	Baxter et al.	
	CPC	2004/0078467	A1 *	4/2004	Grosner	H04L 67/1095 709/226
	<i>G06F 17/30283</i> ; <i>G06F 17/30193</i> ; <i>G06F 17/30997</i> ; <i>G06F 17/30442</i> ; <i>G06F 17/30477</i> ; <i>G06F 17/30545</i> ; <i>G06F 17/30</i> ; <i>G06F 17/07</i> ; <i>G06F 17/30067</i> ; <i>G06F 17/30106</i> ; <i>G06F 17/30306</i> ; <i>G06F 17/30424</i> ; <i>G06F 17/30023</i> ; <i>G06F 17/3007</i> ; <i>G06F 16/24542</i> ; <i>G06F 16/10</i> ; <i>G06F 16/11</i> ; <i>G06F 16/27</i> ; <i>G06F 16/148</i> ; <i>G06F 16/182</i> ; <i>G06F 16/907</i> ; <i>G06F 16/1858</i> ; <i>G06F 16/2453</i> ; <i>G06F 16/2471</i> ; <i>G06F 16/24524</i> ; <i>G06F 16/24532</i> ; <i>G04L 67/1097</i> ; <i>H04L 65/60</i> ; <i>H05K 999/99</i>	2004/0088282	A1	5/2004	Xu et al.	
	See application file for complete search history.	2004/0098415	A1	5/2004	Bone et al.	
		2005/0091287	A1	4/2005	Sedlar	
		2005/0198401	A1	9/2005	Chron et al.	
		2006/0037069	A1	2/2006	Fisher et al.	
		2006/0173751	A1	8/2006	Schwarze et al.	
		2007/0282951	A1	12/2007	Selimis	
		2008/0059746	A1	3/2008	Fisher	
		2008/0172281	A1	7/2008	Probst	
		2008/0320151	A1	12/2008	McCanne et al.	
		2009/0106255	A1	4/2009	Lacapra et al.	
		2009/0222569	A1	9/2009	Frick	
		2009/0254916	A1	10/2009	Bose et al.	
		2010/0036840	A1	2/2010	Pitts	
		2010/0274772	A1	10/2010	Samuels	
(56)	References Cited	2010/0287170	A1	11/2010	Liu et al.	
	U.S. PATENT DOCUMENTS	2011/0137966	A1	6/2011	Srinivasan	
		2011/0179250	A1	7/2011	Matsuzawa	
		2012/0036107	A1	2/2012	Miloushev et al.	
	5,922,030 A	2012/0089470	A1	4/2012	Barnes, Jr.	
	7/1999 Shank	2012/0095952	A1	4/2012	Archambeau et al.	
	6,266,682 B1	2012/0101991	A1	4/2012	Srivasa	
	7/2001 Lamarca et al.	2012/0166483	A1	6/2012	Choudhary	
	6,269,380 B1	2012/0278471	A1	11/2012	Labowicz	
	7/2001 Terry	2013/0036272	A1	2/2013	Nelson	
	6,718,372 B1 *	2013/0262443	A1	10/2013	Leida	
	4/2004 Bober	2014/0136483	A1	5/2014	Chaudhary et al.	
	<i>G06F 17/30203</i> 707/E17.032	2014/0195558	A1	7/2014	Murthy	
		2016/0150019	A1	5/2016	Klinkner	

* cited by examiner

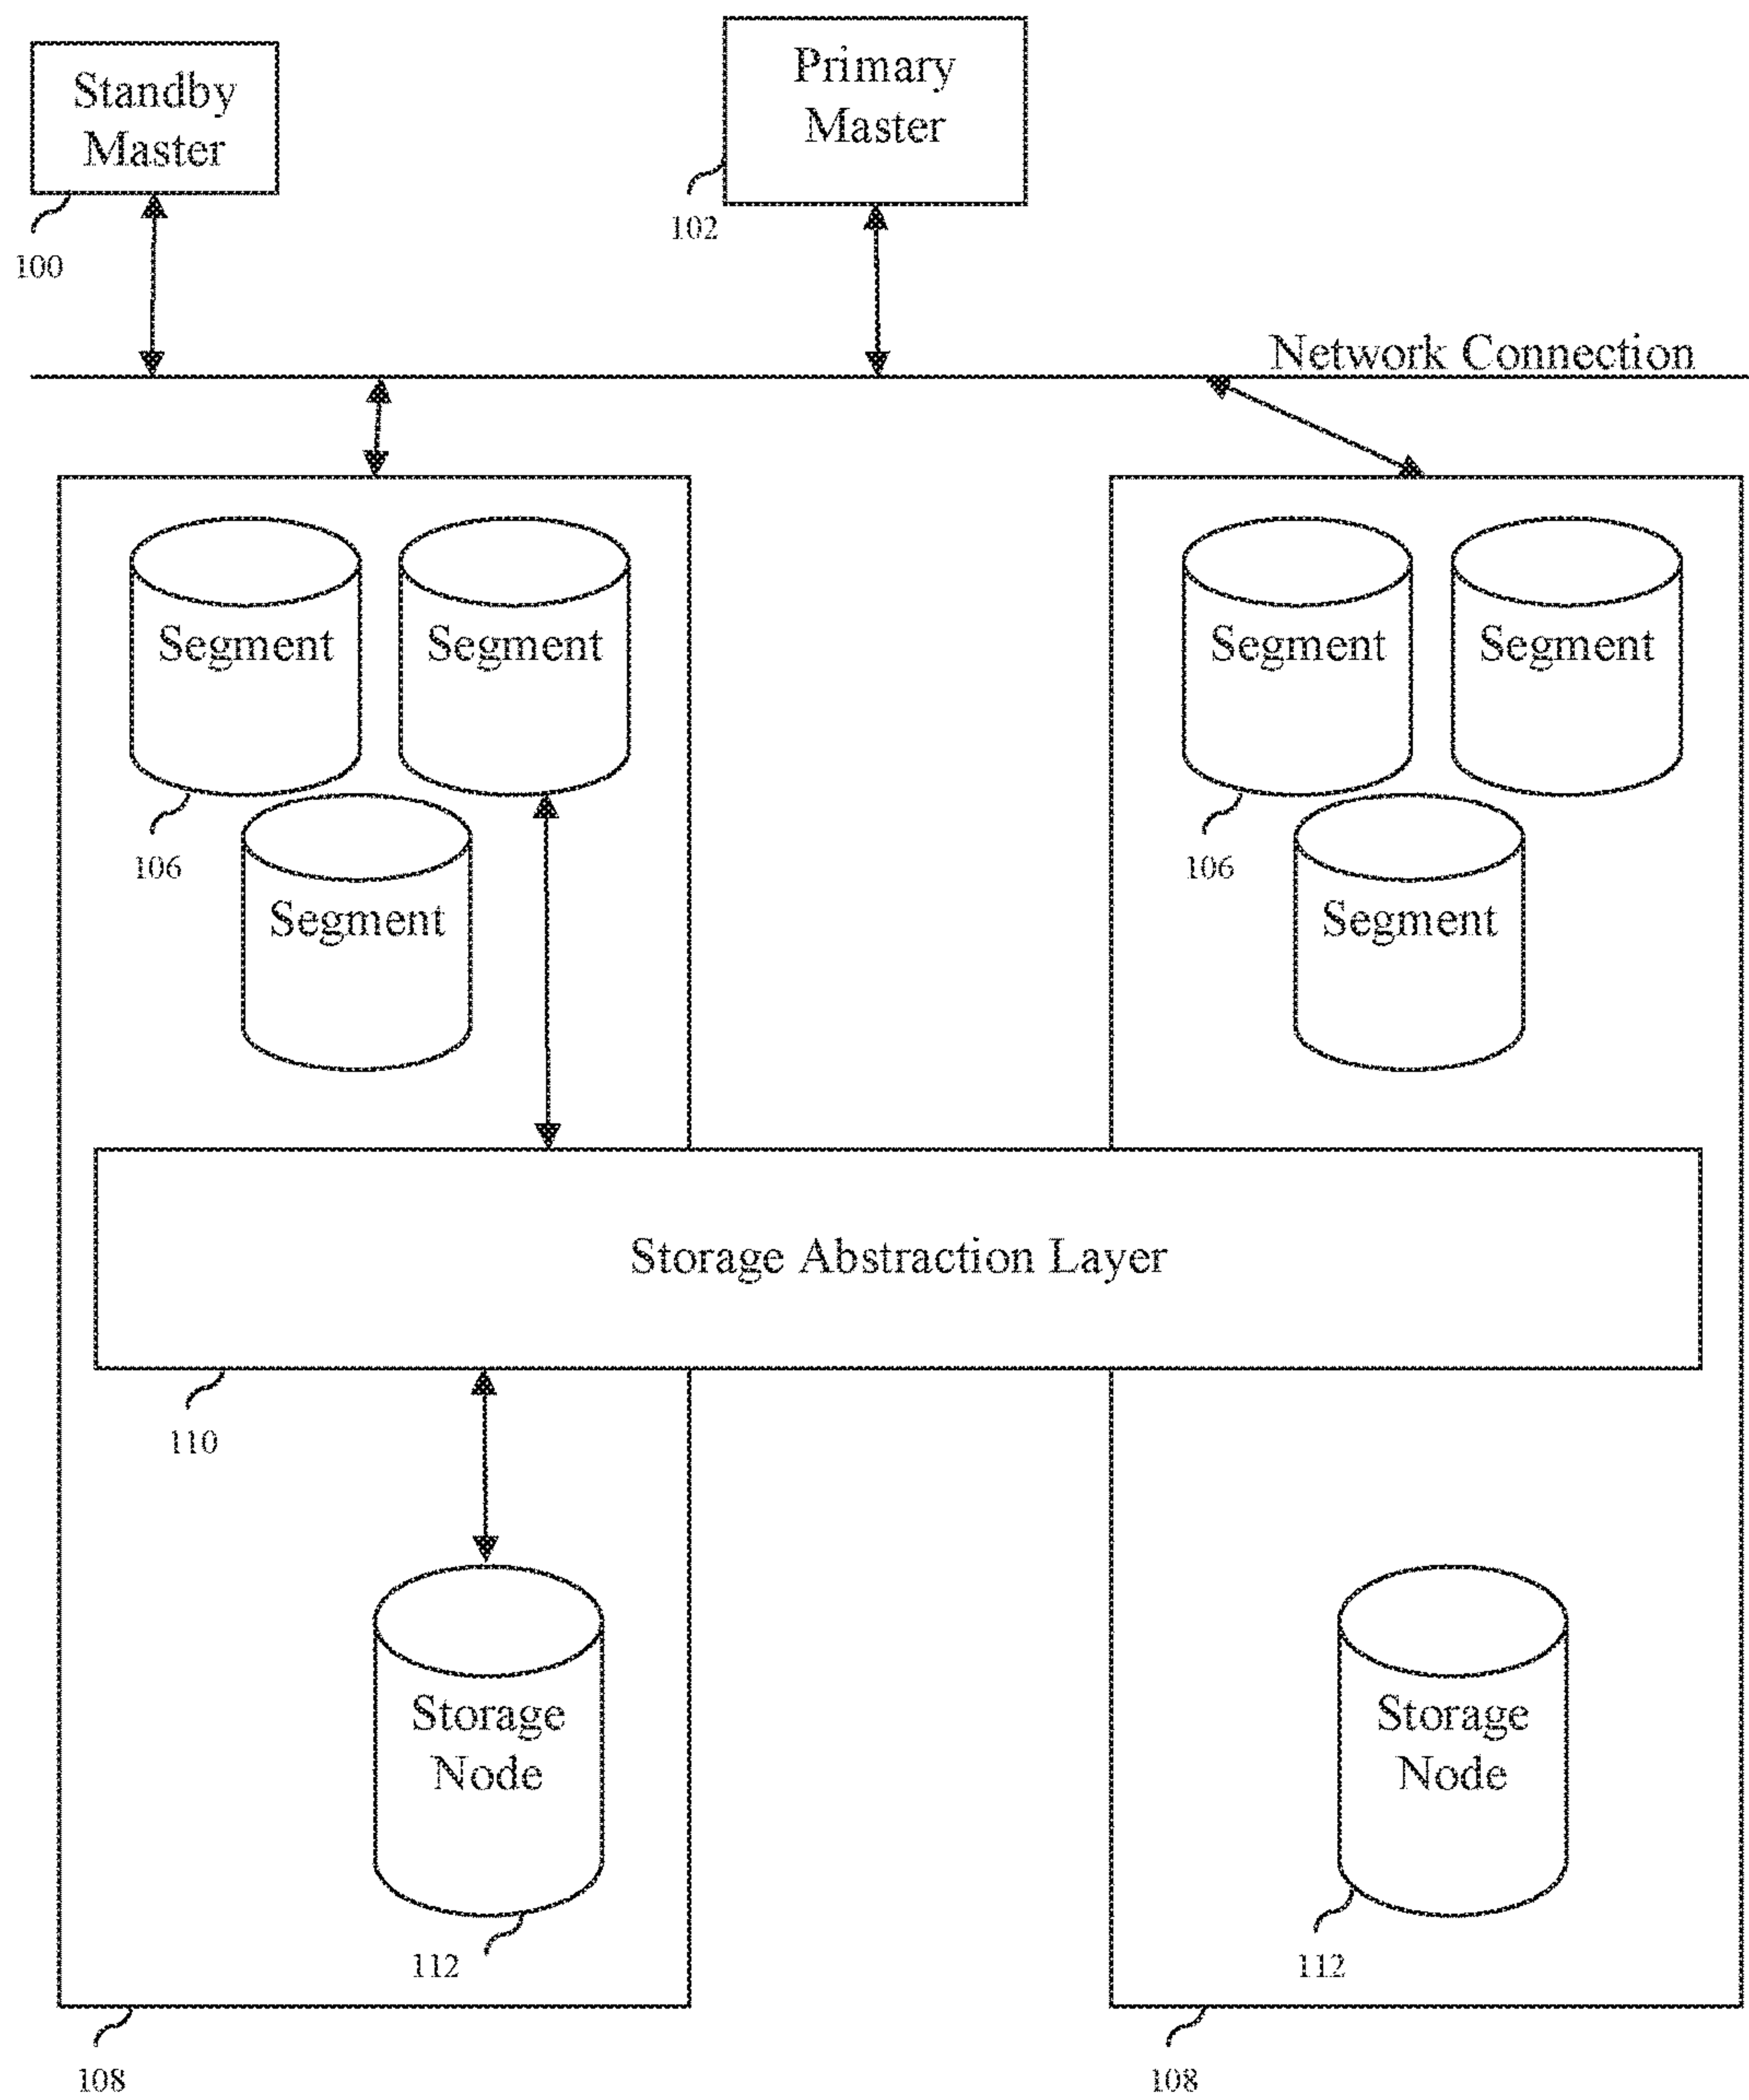


FIG. 1

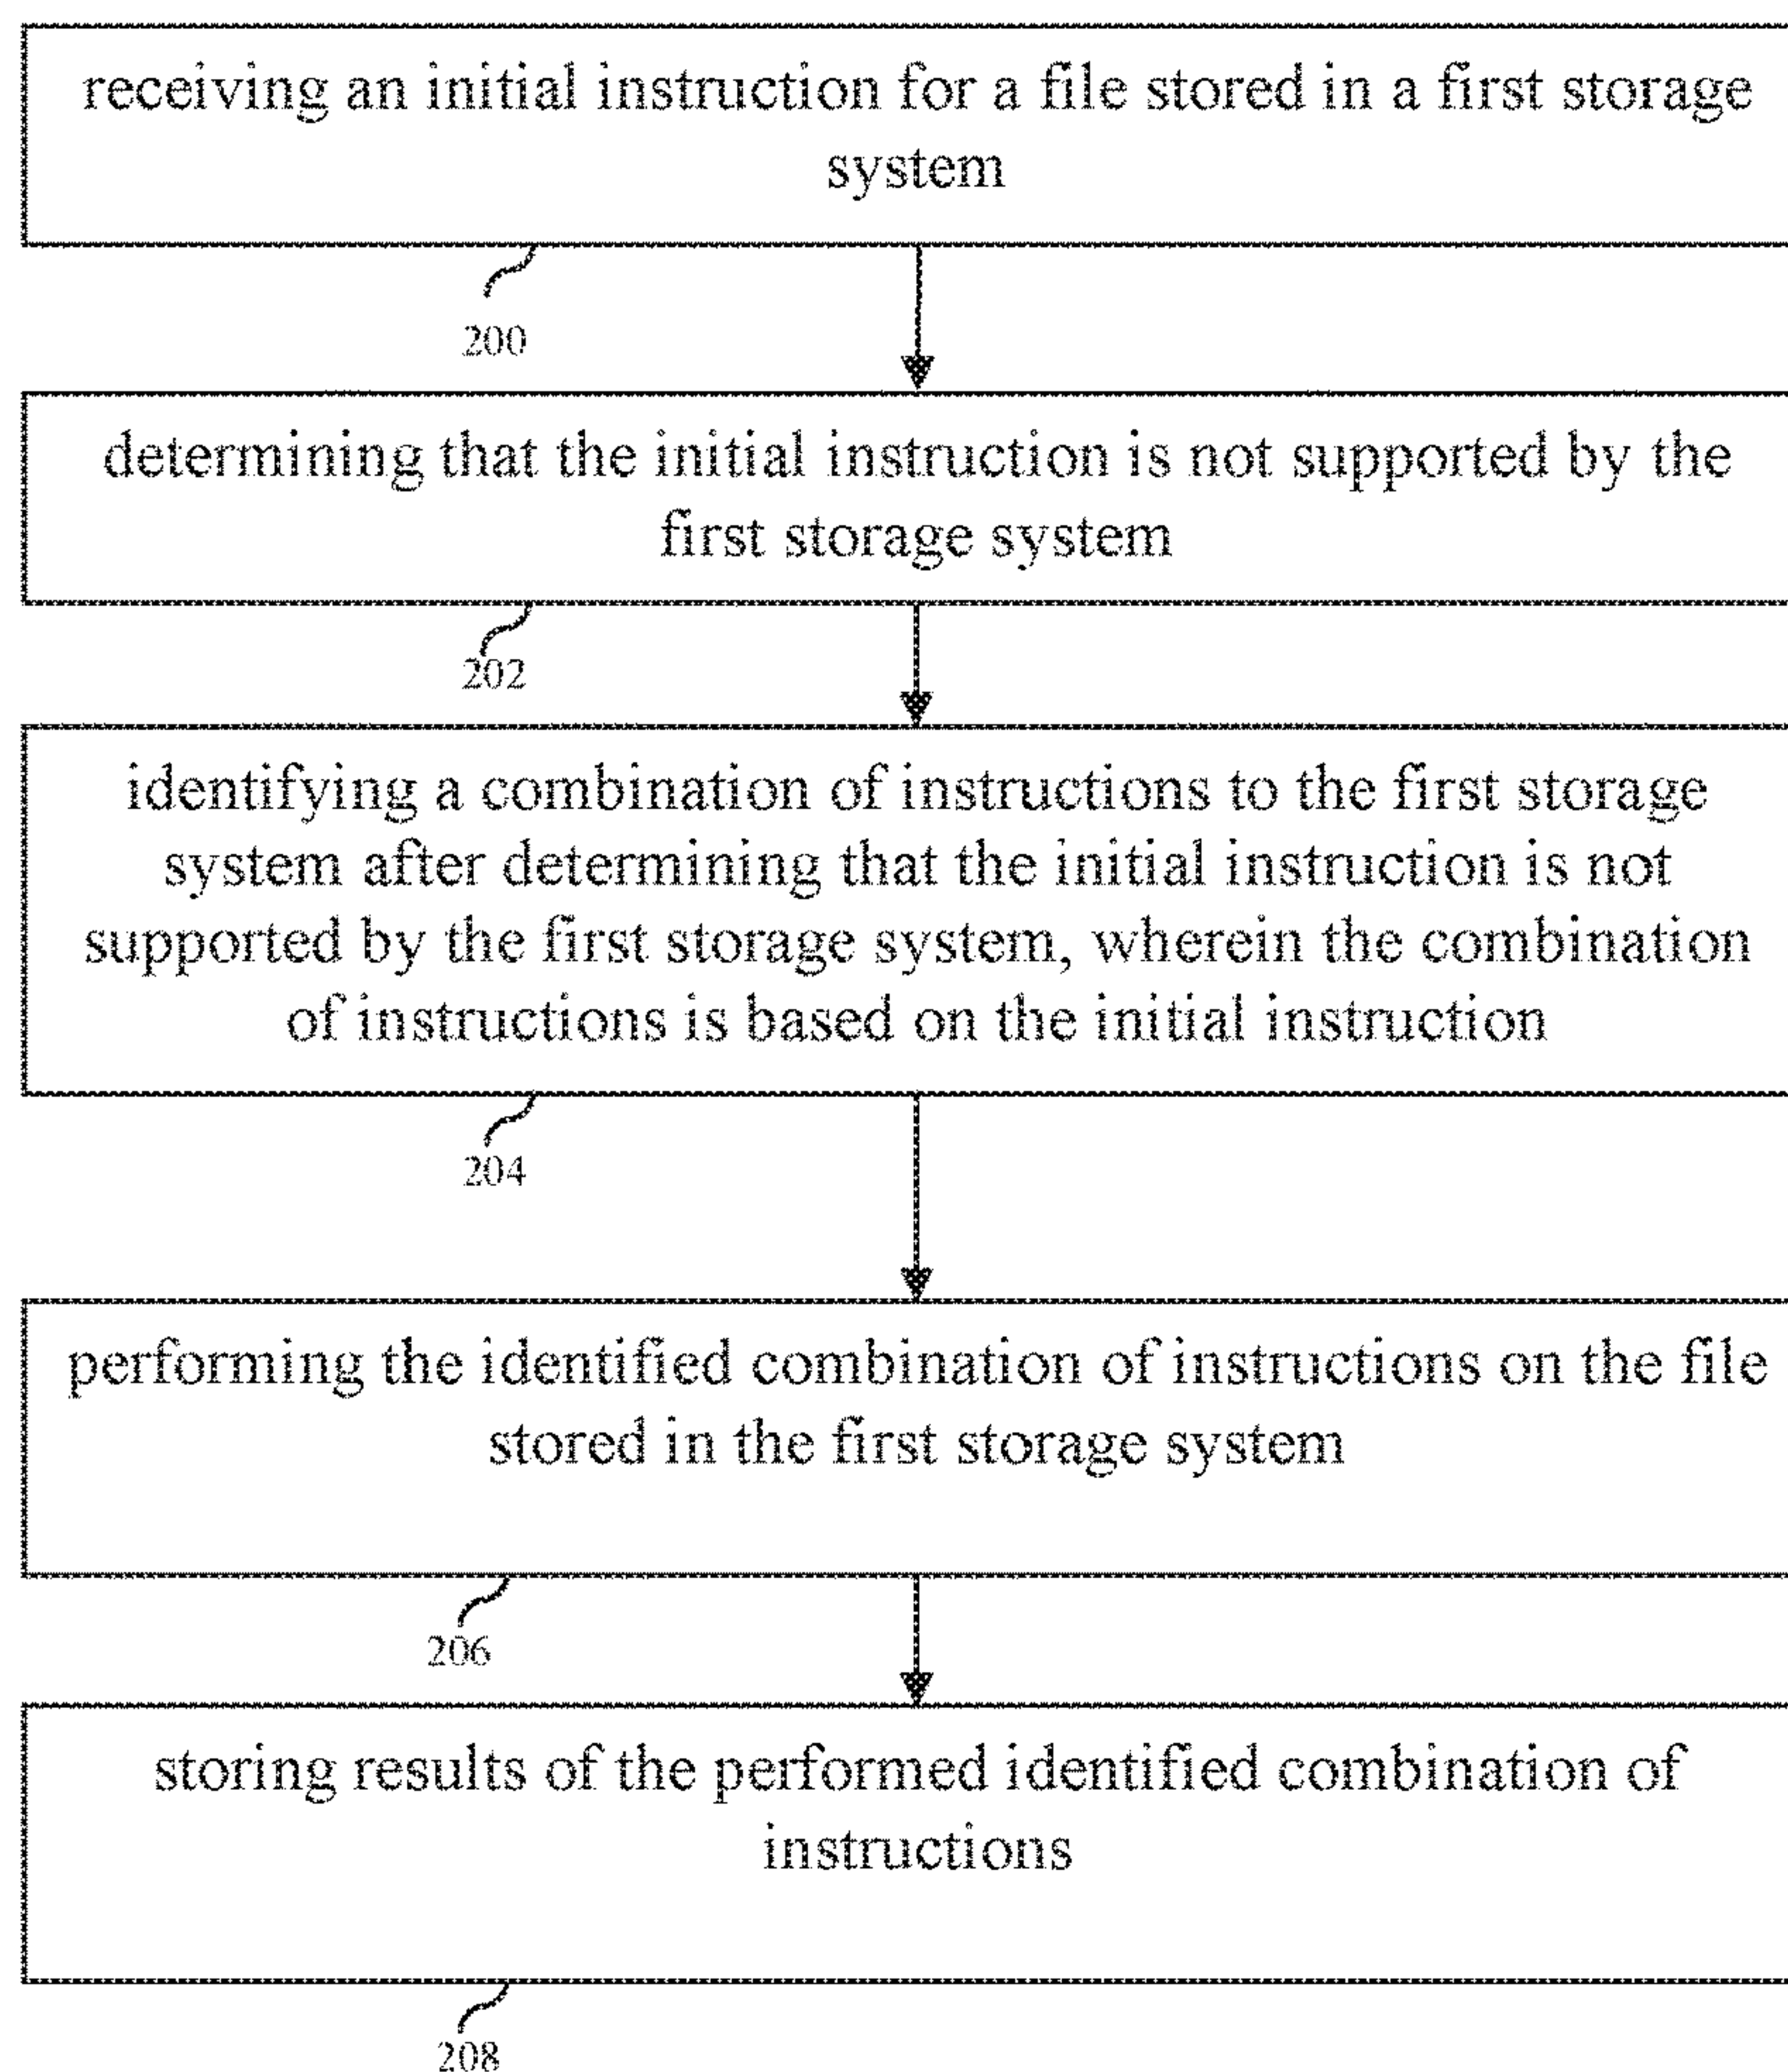


FIG. 2

PLUGGABLE STORAGE SYSTEM FOR DISTRIBUTED FILE SYSTEMS

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation of co-pending U.S. patent application Ser. No. 13/843,067, entitled PLUGGABLE STORAGE SYSTEM FOR DISTRIBUTED FILE SYSTEMS filed Mar. 15, 2013 which is incorporated herein by reference for all purposes, which claims priority to U.S. Provisional Application No. 61/769,043, entitled INTEGRATION OF MASSIVELY PARALLEL PROCESSING WITH A DATA INTENSIVE SOFTWARE FRAMEWORK filed Feb. 25, 2013 which is incorporated herein by reference for all purposes.

FIELD OF THE INVENTION

This invention relates generally to databases, and more particularly to systems and methods for managing datasets in databases.

BACKGROUND OF THE INVENTION

With the large amounts of data generated in recent years, data mining and machine learning are playing an increasingly important role in today's computing environment. For example, businesses may utilize either data mining or machine learning to predict the behavior of users. This predicted behavior may then be used by businesses to determine which plan to proceed with, or how to grow the business.

The data used in data mining and analytics is typically not stored in a uniform data storage system. Many data storage systems utilize different file systems, and those different file systems are typically not compatible with each other. Further, the data may reside in geographically diverse locations.

One conventional method to performing data analytics across different databases includes copying data from one database to a central database, and performing the data analytics on the central database. However, this results in an inefficient use of storage space, and creates issues with data consistency between the two databases.

There is a need, therefore, for an improved method, article of manufacture, and apparatus for managing data.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

FIG. 1 illustrates a database system in accordance with some embodiments.

FIG. 2 is a flowchart of a method to manage data in accordance with some embodiments.

DETAILED DESCRIPTION

A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. While the invention is described in conjunction with such embodiment(s), it should be understood that the invention is not limited to any one embodiment. On the contrary, the scope of the invention is limited only by the claims and the

invention encompasses numerous alternatives, modifications, and equivalents. For the purpose of example, numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention. These details are provided for the purpose of example, and the present invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the present invention is not unnecessarily obscured.

It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device, a method, or a computer readable medium such as a computer readable storage medium or a computer network wherein computer program instructions are sent over optical or electronic communication links. Applications may take the form of software executing on a general purpose computer or be hardwired or hard coded in hardware. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention.

An embodiment of the invention will be described with reference to a data storage system in the form of a storage system configured to store files, but it should be understood that the principles of the invention are not limited to this configuration. Rather, they are applicable to any system capable of storing and handling various types of objects, in analog, digital, or other form. Although terms such as document, file, object, etc. may be used by way of example, the principles of the invention are not limited to any particular form of representing and storing data or other information; rather, they are equally applicable to any object capable of representing information.

FIG. 1 illustrates a database system in accordance with some embodiments. Primary Master **102** accepts queries, plans queries, and dispatches queries to Segments **106** for execution. Primary Master **102** also collects results from Segments **106**. Segments **106** are each a compute unit of the database system. Rack **108** can store multiple segments as hosts (hosts not shown in FIG. 1). Standby Master **100** is a warm backup of the primary host. Primary Master **102** and Standby Master **100** communicate to Rack **108** via the network connection. The tables of the database system are stored in Storage Nodes **112**. Access, including read and write, are done through Storage Abstraction Layer **110**. Primary Master **102**, in some embodiments, may include a meta store (not shown in FIG. 1).

The meta store includes information about the different files systems in Storage Nodes **112**, such as API information of the file system's interface, and different attributes and metadata of the file system. The meta store also includes information on the binary location of the Storage Abstraction Layer **110**. As new file systems are added to the database system, the new file systems are registered (e.g. provided API information, other attributes, etc.) with the meta store. Once the new file systems are added, instances of that file system may be created to store data objects, such as databases and tables.

The storage nodes may be different file systems. For example, one storage node may be Hadoop File System (HDFS), while another storage node may be NFS. Having multiple file systems presents some challenges. One chal-

lenge is that file systems do not support all the same commands. The Storage Abstraction Layer helps address this challenge.

In some embodiments, the Storage Abstraction Layer selects a file system instance. A file system instance means a physical storage system for a specific file system. As discussed above, there may be several different file systems, and several different instances. The instances may be of the same file system, or they may be of different file systems. Different file systems may have different semantics or different performance characteristics. For example, some file systems allow you to update data, while other file systems only let you append data. The Storage Abstraction Layer chooses a file system based on the file system's attributes.

For example, in some embodiments, if a user wanted to modify or update a file that is stored on an underlying storage system which does not support file modification, the Storage Abstraction Layer may recognize the update command and move the file from the underlying storage system to another which does support file modification. The move may be temporary (e.g. move the file back after the user is finished with the file), or the move may be permanent.

In some embodiments, the Storage Abstraction Layer may choose to store a data object in a file system that does not allow updating. This may be preferable in cases where the data object is only read and never modified, and the file system is efficient for retrieving data. Thus, the Storage Abstraction Layer may take into account the usage statistics of the data object to determine what file system to use to store the data.

In some embodiments, the Storage Abstraction Layer may perform semantic adaptation. This may be preferable when the underlying file system may not be able to communicate directly with segments. This may occur when the interface the Storage Abstraction Layer exposes to the segment execution engine does not match with the semantics of the underlying file system. Other examples include instances where the functionality required by the segments is not supported by the underlying file system.

For example, a user may wish to truncate a file. However, the file may be stored on a segment where the underlying storage does not allow truncating files. The user is not aware of this because the user is not aware of where the files are physically stored. Typically, without a Storage Abstraction Layer, the underlying file system would not be able to understand the truncate command.

One example of semantic adaptation includes adapting the truncate command. Suppose that a segment requires a piece of data to be truncated. However, the underlying file system does not support the truncate functionality. The Storage Abstraction Layer may be able to put various commands together to mimic a truncate command. Since the Storage Abstraction Layer has access to the metadata of the file system stored in the meta store, it knows what commands are allowed in the file system, as well as how to access the file system via APIs. Suppose that the file to be truncated is File A, and File A consists of 20 bytes. The segment wants the last 10 bytes to be deleted. With this requirement, the Storage Abstraction Layer may employ semantic adaptation to complete the truncation even when the underlying file system does not support a truncate command. In some embodiments, the Storage Abstraction Layer may first copy the first 10 bytes of File A to a temporary file, called File B. Then, the original File A is deleted, leaving only the temporary File B. After the original File A is deleted, the temporary File B is renamed to File A. File A is now only half of the original File A. In other words, File A has been

truncated, even though the underlying file system did not support truncation. The Storage Abstraction Layer, by understanding how to access the underlying file system via the meta store, sent a series of commands to mimic a truncate. This series of commands may be stored in the meta store so that future truncate requests may make use of it.

Another example of semantic adaptation includes a file update command. As mentioned above, some file systems do not allow for updating a file. Suppose a segment required that a file be updated. However, the file is stored in a file system that does not allow files to be updated. In some embodiments, the Storage Abstraction Layer may record the modifications in separate file as a new version. For example, if File A was to be modified, the separate file may be called File A_ver2. The segment (or user) will see that changes are being made to File A, but in fact, File A remains unchanged and the changes are being stored in File A_ver2. After the segment is finished modifying or updating the file, there may be two files stored—one is File A, and the other is File A_ver2. When subsequent users want to access File A, the Storage Abstraction Layer may cause the two files to be merged. With File A merged with File A_ver2 and called File A, the new File A will include all the changes made by the previous user. In other words, File A has been modified, even though the underlying file system did not support updating.

With the Storage Abstraction Layer, many different file systems may be supported. New and different storage systems with different file systems may be “plugged” into the database, without affecting the ability for the database to run its queries or jobs, as long as the meta store is updated with information about the new file system, such as its APIs.

FIG. 2 illustrates a method to manage data in accordance with some embodiments. In step 200, an initial instruction for a file stored in a first storage system is received. In step 202, it is determined that the initial instruction is not supported by the first storage system. In step 204, a combination of instructions to the first storage system is identified after determining that the initial instruction is not supported by the first storage system, wherein the combination of instructions is based on the initial instruction. In step 206, the identified combination of instructions is performed on the file stored in the first storage system. In step 208, results of the performed identified combination of instructions are stored.

For the sake of clarity, the processes and methods herein have been illustrated with a specific flow, but it should be understood that other sequences may be possible and that some may be performed in parallel, without departing from the spirit of the invention. Further, though the techniques herein teach creating one SwR sample in parallel, those with ordinary skill in the art will readily appreciate that the techniques are easily extendable to generate many SwR samples. Additionally, steps may be subdivided or combined. As disclosed herein, software written in accordance with the present invention may be stored in some form of computer-readable medium, such as memory or CD-ROM, or transmitted over a network, and executed by a processor.

All references cited herein are intended to be incorporated by reference. Although the present invention has been described above in terms of specific embodiments, it is anticipated that alterations and modifications to this invention will no doubt become apparent to those skilled in the art and may be practiced within the scope and equivalents of the appended claims. More than one computer may be used, such as by using multiple computers in a parallel or load-sharing arrangement or distributing tasks across multiple computers such that, as a whole, they perform the functions

5

of the components identified herein; i.e. they take the place of a single computer. Various functions described above may be performed by a single process or groups of processes, on a single computer or distributed over several computers. Processes may invoke other processes to handle certain tasks. A single storage device may be used, or several may be used to take the place of a single storage device. The disclosed embodiments are illustrative and not restrictive, and the invention is not to be limited to the details given herein. There are many alternative ways of implementing the invention. It is therefore intended that the disclosure and following claims be interpreted as covering all such alterations and modifications as fall within the true spirit and scope of the invention.

What is claimed is:

1. A method for managing data, comprising:
 receiving, by an interfacing module between a segment execution engine and a first storage system, an initial instruction for a file stored in the first storage system;
 determining, by the interfacing module, that the initial instruction is not supported by the first storage system;
 identifying, by the interfacing module, a combination of instructions to the first storage system after determining that the initial instruction is not supported by the first storage system, wherein the combination of instructions is based on the initial instruction;
 performing the identified combination of instructions on the file stored in the first storage system; and
 storing results of the performed identified combination of instructions.

2. The method of claim 1, wherein the combination of instructions is determined based at least in part on metadata associated with the first storage system, and the metadata associated with the first storage system is stored in a meta store that includes information respectively associated with one or more storage systems.

3. The method of claim 1, wherein the combination of instructions is determined based at least in part on metadata associated with the first storage system, and the metadata associated with the first storage system indicates a mechanism for accessing a file system corresponding to the first storage system.

4. The method of claim 1, wherein the combination of instructions comprise one or more instructions that are determined based at least in part on metadata associated with the first storage system.

5. The method of claim 1, wherein the combination of instructions comprise one or more instructions that are determined based at least in part on the initial instruction.

6. The method of claim 1, further comprising:
 selecting a second storage system to which the file is to be moved;
 storing the file on the second storage system; and
 performing a series of actions in connection with carrying out the collective result.

7. The method of claim 6, wherein the combination of instructions comprise one or more instructions, the one or more instructions include one or more of the selecting of the second storage system, the storing of the file on the second storage system, and the performing of the series of actions in connection with carrying out the collective result.

8. The method of claim 6, wherein the second storage system is selected based at least in part on the second storage system having a second file system.

9. The method of claim 6, further comprising:
 in response to performing the series of actions, moving the file to the first storage system.

6

10. The method of claim 1, wherein the initial instruction includes a truncate instruction.

11. The method of claim 10, wherein the combination of instructions comprise one or more instructions, the one or more instructions include copy a first portion of the file, delete the file, and rename the first portion of the file.

12. The method of claim 1, wherein the initial instruction includes an update instruction.

13. A system for managing data, comprising a storage device and at least one processor configured to:

receive, by an interfacing module between a segment execution engine and a first storage system, an initial instruction associated with performing a first action in relation to a file stored in the first storage system, wherein the file is stored in the first storage system based at least in part on file system functionality associated with a first file system associated with the first storage system and one or more usage statistics associated with the file;

determine, by the interfacing module, that the initial instruction associated with performing the first action is not supported by the first storage system based at least in part on metadata associated with the first storage system;

in response to determining that the initial instruction associated with performing the first action is not supported by the first storage system, identify, by the interfacing module, one or more instructions that have a collective result that is equivalent to a result of the first action, wherein to identify the combination of instructions includes to store an update in a second file and merging the second file with the file after the update is complete;

perform the identified one or more instructions in connection with the file stored in the first storage system; and

store results of the performed identified one or more instructions.

14. The system of claim 13, wherein the combination of instructions is determined based at least in part on metadata associated with the first storage system, and the metadata associated with the first storage system is stored in a meta store that includes information respectively associated with one or more storage systems.

15. The system of claim 13, wherein the combination of instructions is determined based at least in part on metadata associated with the first storage system, and the metadata associated with the first storage system indicates a mechanism for accessing a file system corresponding to the first storage system.

16. The system of claim 13, wherein the combination of instructions comprise one or more instructions that are determined based at least in part on metadata associated with the first storage system.

17. The system of claim 13, wherein the combination of instructions comprise one or more instructions that are determined based at least in part on the initial instruction.

18. The system of claim 13, wherein the at least one processor is further configured to:

select a second storage system to which the file is to be moved;

store the file on the second storage system; and
 perform a series of actions in connection with carrying out the collective result.

19. The system of claim 18, wherein the combination of instructions comprise one or more instructions that cause the processor to perform one or more of the selecting of the

7

second storage system, the storing of the file on the second storage system, and the performing of the series of actions in connection with carrying out the collective result.

20. A computer program product for processing data, comprising a non-transitory computer readable medium having program instructions embodied therein for:

receiving, by an interfacing module between a segment execution engine and a first storage system, an initial instruction associated with performing a first action in relation to a file stored in a first storage system, wherein the file is stored in the first storage system based at least in part on file system functionality associated with a first file system associated with the first storage system and one or more usage statistics associated with the file; determining, by the interfacing module, that the initial instruction associated with performing the first action is

8

not supported by the first storage system based at least in part on metadata associated with the first storage system;

in response to determining that the initial instruction associated with performing the first action is not supported by the first storage system, identifying, by the interfacing module, one or more instructions that have a collective result that is equivalent to a result of the first action, wherein the identifying of the combination of instructions includes storing an update in a second file and merging the second file with the file after the update is complete;

performing the identified one or more instructions in connection with the file stored in the first storage system; and

storing results of the performed identified one or more instructions.

* * * * *