STEPHEN M. LOBBIN
California Bar No. 181195
sml@smlavvocati.com
SML AVVOCATI P.C.
888 Prospect Street, Suite 200
San Diego, CA 92037
Telephone: 949.636.1391

CABRACH J. CONNOR
Texas Bar No. 24036390 (*pro hac vice* pending)
cab@connorkudlaclee.com
KEVIN S. KUDLAC
Texas Bar No. 00790089 (*pro hac vice* pending)
kevin@connorkudlaclee.com
JENNIFER TATUM LEE
Texas Bar No. 24046950 (*pro hac vice* pending)
jennifer@connorkudlaclee.com
CONNOR KUDLAC LEE PLLC
609 Castle Ridge Road, Suite 450
Austin, TX 78746
Telephone: 512.777.1254
Facsimile: 888.387.1134

Attorneys for Plaintiff
MOBILE NETWORKING SOLUTIONS, LLC

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

MOBILE NETWORKING SOLUTIONS, LLC, Plaintiff, vs. ALIBABA GROUP HOLDING LTD., ALIBABA GROUP (U.S.) INC., and ALIBABA CLOUD US LLC, Defendants.	CASE NO. 3:19-cv-00217-KAW_____ **COMPLAINT FOR PATENT INFRINGEMENT** **DEMAND FOR JURY TRIAL**

Mobile Networking Solutions, LLC ("MNS") files this Complaint for Patent Infringement against Alibaba Group Holding Limited, Alibaba Group (U.S.) Inc., and Alibaba Cloud US LLC ("AliCloud") for infringement of U.S. Patents Nos. 7,543,177 and 7,958,388 relating to large-scale data storage, processing, and management.

PARTIES

1. MNS is a limited liability company organized and existing under the laws of the State of Texas with its principal place of business at 1400 Preston Road, Suite 483, Plano, Texas 75093.

2. Upon information and belief, Defendant Alibaba Group Holding Ltd. is a Cayman Islands holding company established under the Companies Law of the Cayman Islands.

3. Alibaba Group Holding Ltd. is headquartered at 969 West Wen Yi Road, Yu Hang District, Hangzhou 311121, People's Republic of China (Telephone: 86-571-8502-2088) and may be served through its registered agent for process: Corporation Service Company, 1180 Avenue of the Americas, Suite 210, New York, New York 100376.

4. Alibaba Group (U.S.) Inc., formerly known as Nimbus Development, Inc.,[1] is a Delaware corporation with its principal office at 400 South El Camino Real, Suite 400, San Mateo, California 94402. Alibaba Group (U.S.) may be served with process through its registered agent, Hong Tang, at its principal office.

5. Alibaba Cloud US LLC is a Delaware limited liability company and may be served with process through its registered agent, National Registered Agents, Inc., 150 Greentree Drive, Suite 101, Dover, Delaware 19904.

6. Alibaba Cloud is Alibaba Group's cloud computing arm.[2] Alibaba Cloud offers a suite of global cloud computing services including elastic computing, database, storage, network virtualization services, large scale computing, security, management, and application services, big data analytics, a machine learning platform, and IoT services to support participants of Alibaba Group's online and mobile commerce ecosystem, including sellers, and other third-party customers and businesses." Alibaba Cloud is also known as Aliyun or "AliCloud" and is headquartered in Singapore.

---

[1] Nimbus Development notified the California Secretary of State of this name change by filing an Amended Statement by Foreign Corporation on July 31, 2017.

[2] https://www.alibabacloud.com/press-room/alibaba-to-open-second-silicon-valley-data-center-to-meet?spm=a2c5t.10695662.1996646101.searchclickresult.7d0041092iUd3d.

7. Alibaba Cloud operates multiple data centers in the United States. The US East 1 data center in Virginia and US West 1 data centers in Northern California.

Region	City	Region ID	Number of zones
Hong Kong	Hong Kong	cn-hongkong	2
Asia Pacific SE 1	Singapore	ap-southeast-1	3
Asia Pacific SE 2	Sydney	ap-southeast-2	2
Asia Pacific SE 3	Kuala Lumpur	ap-southeast-3	2
Asia Pacific SE 5	Jakarta	ap-southeast-5	1
Asia Pacific SOU 1	Mumbai	ap-south-1	2
Asia Pacific NE 1	Tokyo	ap-northeast-1	1
US West 1	Silicon Valley	us-west-1	2
US East 1	Virginia	us-east-1	2

8. Alibaba Cloud regions are physical locations (data centers) . Zones are the physical areas with independent power grids and networks in one region.[3]

9. Alibaba Cloud's U.S. data centers are intended to serve business operations in America and Continental Europe:

- North America and South America
  - **US West 1**

    The data center in US West 1 is located in Silicon Valley. It is directly connected to the backbone networks of multiple American operators through BGP lines. In addition to the United States, this data center extends its reach to South America and Continental Europe. If you have business operation in America and Continental Europe, select this region.
  - **US East 1**

    The data center in US East 1 is located in Virginia of the United States. If you have business operation in America and Continental Europe, select this region.

https://www.alibabacloud.com/help/doc-detail/40654.htm.

10. On July 1, 2017, Alibaba Cloud US LLC replaced Nimbus Development, Inc. as the provider of cloud services to users residing in or whose billing address is located in the United States.[4]

11. Purchasers of Alibaba Cloud products in the United States and purchasers with a billing address in the Unites States contract with Alibaba Cloud US LLC.[5]

[3] https://www.alibabacloud.com/help/doc-detail/40654.htm.
[4] https://www.alibabacloud.com/notice/Legal-Notice?spm=0.lp622.0.0.63cf5a47S61Y7j.
[5] https://www.alibabacloud.com/help/faq-detail/42416.htm?spm=a2145.520000001.0.0.7c8152cbAdGjzS.

JURISDICTION AND VENUE

12. This is an action under the patent laws of the United States, 35 U.S.C. §§ 1, *et seq*. and namely §§ 271, 281, and 284-285, for infringement by Alibaba and AliCloud of claims of U.S. Patent Nos. 7,543,177 (the "'177 Patent") and 7,958,388 (the "'388 Patent") (collectively, the "Patents-in-Suit").

13. This Court has subject-matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

14. Alibaba Group, Alibaba Group (U.S.), and Alibaba Cloud US (collectively referred to as "Alibaba" or the "Alibaba Defendants") are subject to general and specific personal jurisdiction of this Court based upon their regularly conducted business in California and in this judicial district giving rise to this action and have established minimum contacts with this forum such that the exercise of jurisdiction over the Alibaba Defendants would not offend traditional notions of fair play and substantial justice.

15. The Alibaba Defendants, directly and through subsidiary business units have committed and continue to commit acts of infringement in this district pursuant to 35 U.S.C. § 271(a) by making, using, selling, offering to sell, testing, deploying, and exercising control and obtaining beneficial use in this district of products and services that infringe the asserted MNS patents.

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

17. Venue is proper in this district as to AliCloud and Alibaba Group (U.S.) Inc. pursuant to 28 U.S.C. § 1400(b) as AliCloud and Alibaba Group (U.S.) Inc. maintain regular and established places of business in this judicial district (offices and data centers) and have committed acts of infringement in this district.

18. Alibaba Group Holding Limited is a corporation organized under the laws of the Cayman Islands. Venue is proper in this district as to Alibaba Group Holding Limited pursuant to 28 U.S.C. § 1400(b) and 28 U.S.C. § 1391. Alibaba Group Holding Limited has committed acts of infringement in this district. Alibaba Group Holding Limited has a

regular and established place of business in this district at 400 South El Camino Real, Suite 400, San Mateo, California 94402. *See* https://www.alibabagroup.com/en/contact/offices (last visited, November 26, 2018).

THE MNS PATENTS

19. MNS is the owner by assignment of all right, title, and interest in and to U.S. Patent Nos. 7,543,177 and 7,958,388 (the "Asserted Patents"), both titled, "Methods and Systems for a Storage System."

20. A true and correct copy of the '177 patent is attached as Exhibit A.

21. A true and correct copy of the '388 Patent is attached as Exhibit B.

22. MNS possesses all rights of recovery under the Asserted Patents.

23. The Asserted Patents issued from continuations of Application No. 10/284,199 filed on October 31, 2002.

24. The U.S. Patent Office issued the '177 Patent on June 2, 2009, after a full examination based upon an application filed by inventors Melvin James Bullen, Steven Louis Dodd, William Thomas Lynch, and David James Herbison.

25. The Examiner stated the following reasons for allowing the claimed subject matter of the '177 Patent:

> Regarding claim 1, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, a switch controller that executes software, including a routing algorithm and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm.
>
> Regarding claim 26, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, a management system determining a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface and the management system removing from service the memory section from which the fault message was received by changing the routing algorithm.
>
> Regarding claim 40, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, programmable means for

> switching data being transmitted between the means for storing and one or more interfaces based on a routing algorithm and means for receiving the fault message, removing from service the means for storing from which the fault message was received by changing the routing algorithm.

26. The U.S. Patent Office issued the '388 Patent on June 7, 2011, after a full examination based upon an application filed by the same inventors.

27. The Examiner stated the following reasons for allowing the claimed subject matter of the '388 Patent: "the prior art does not teach or reasonably suggest providing, by the management system, the routing algorithm to the switch controller and determining, by the management system in response to the detecting, a new routing algorithm that redirects data for the memory device to a replacement memory device; and providing the new routing algorithm to the switch controller."

28. The Abstract of the Asserted Patents describes the claimed subject matter as being directed to "[a] storage system that may include one or more memory sections, one or more switches, and a management system . . . [t]he memory sections include memory devices and a section controller capable of detecting faults with the memory section and transmitting messages to the management system regarding detected faults. The storage system may include a management system capable of receiving fault messages from the section controllers and removing from service the faulty memory sections . . . [a]dditionally, the management system may determine routing algorithms for the one or more switches."

29. Figure 6 in the specification of the Asserted Patents is a functional diagram exemplifying the claimed subject matter:

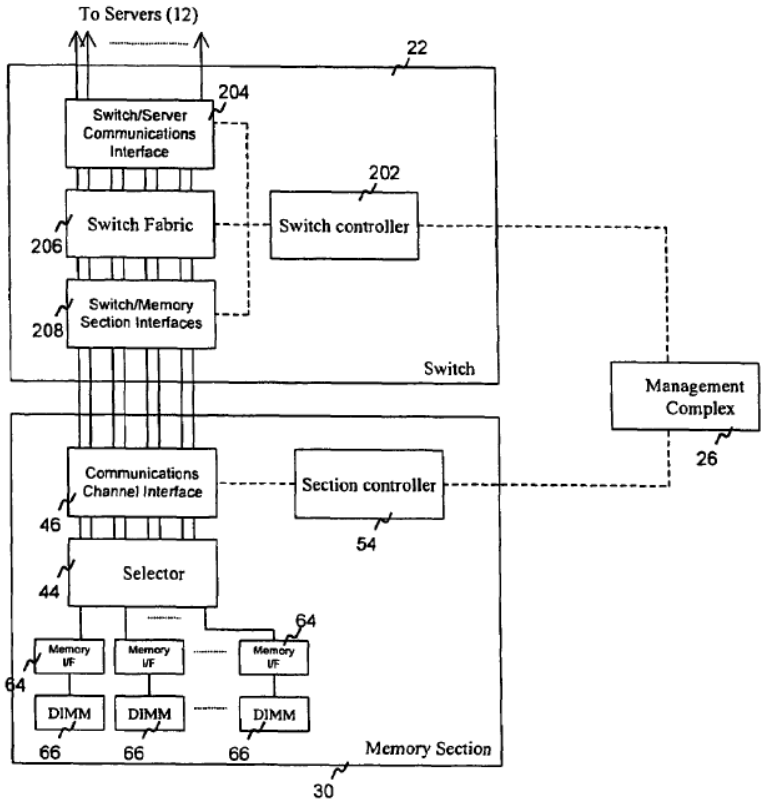

30. The inventors recognized and noted in the specification that large-scale storage systems suffered from problems in throughput for high-volume, real-time applications.

31. In operation, the switches, memory sections, and management system of the Asserted Patents receive fault messages from the memory section controllers and remove from service the memory section from which the fault message was received, and the management system may further determine an algorithm for use by a switch fabric in interconnecting the memory sections and external device interfaces and instruct the switch to executed the determined algorithm. '177 Patent at 2:21-34.

32. Those of skill in the art at the time of the inventions claimed in the Asserted Patent would recognize that the claimed subject matter addresses performance limitations inherent in disk storage technologies such as input/output bottlenecks and improves

network operations in the event of signal and/or equipment failure.

33. The claimed subject matter of the Asserted Patents is particularly applicable to improve the operation of parallel processing technologies in big-data distributed storage systems such as the Hadoop Distributed File System (HDFS).

Hadoop Distributed File System

34. The Hadoop Distributed File System (HDFS) is used for storage and processing of large data files across a cluster of storage hardware.

35. According to AliCloud, Hadoop comprises three main components: HDFS (Hadoop Distributed File System, a "distributed, reliable, scalable and highly fault-tolerant file system for data storage), MapReduce ("a programming model that is designed for large volumes of data in parallel by dividing the work into a set of independent tasks" that "distributes work across hundreds or thousands of servers in a Hadoop cluster"), and YARN (Yet Another Resource Negotiator that is the "resource management layer of Hadoop . . . responsible for managing computing resources in clusters and using them for scheduling users' applications"). https://www.alibabacloud.com/blog/how-to-setup-hadoop-cluster-ubuntu-16-04_593808.

36. An HDFS instance may consist of hundreds or even thousands of servers (DataNodes) that each store part of a large data file.

37. HDFS features high fault tolerance and automatic fault recovery making it suitable for deployment on commodity hardware.

38. Operational advantages of HDFS include efficient processing by executing application instructions near the subject data. HDFS's cluster design and input/output pathing minimizes network congestion and increases throughput.

39. HDFS handles big data, typically 10-100GB or more with diverse data types including structured and unstructured data, economically distributing the computational load across multiple DataNodes.

40. HDFS DataNodes are a cluster of computers capable of executing the workload components such as storing HDFS data blocks and performing block replication.

41. Distributing the computing load across DataNodes requires multiple servers having access to the data, and HDFS meets this need by ensuring that the entire calculation process does not terminate when an error occurs within a HDFS cluster.

42. The NameNode is responsible for keeping track of file system metadata including a list of blocks in an HDFS file and a list of DataNodes.

43. AliCloud publishes the following pictorial representation of the HDFS architecture:

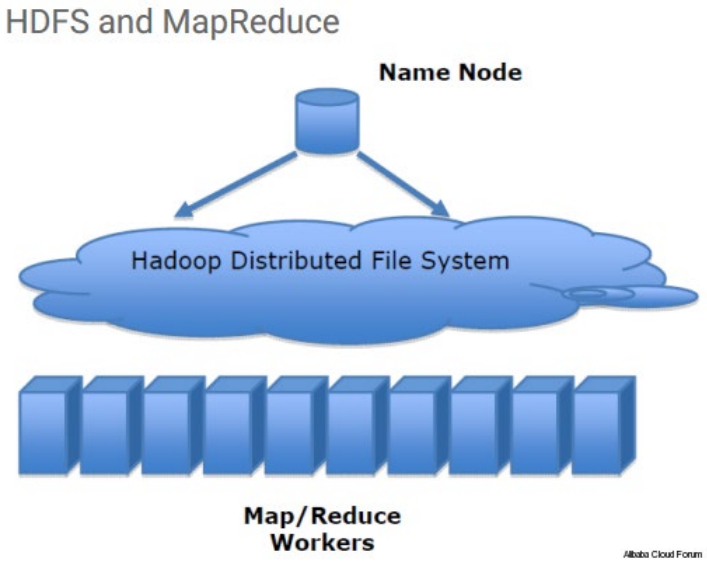

44. An article published by AliCloud focusing on the components of the Hadoop system and functions of the various components describes MapReduce as the "core of Hadoop." https://www.alibabacloud.com/forum/read-896.

45. AliCloud describes how MapReduce decomposes the processing task of a large data set query for processing on multiple running nodes because one server "won't be enough to solve the problem . . . the advantage of distributed computing is apparent."

46. According to AliCloud, "[t]he combination of HDFS and MapReduce is powerful."

47. AliCloud explains:

> During big data processing, the entire calculation process does not terminate when an error occurs in a server in the Hadoop cluster. At the same time, HDFS guarantees data redundancy in the event of a fault or error across the cluster. The result is written to a node in the HDFS when the calculation is

> complete. HDFS has no strict requirements for stored data formats, and data can be unstructured or in other categories. In contrast, relational databases require data to be structured and the architecture defined before storing the data.

https://www.alibabacloud.com/forum/read-896.

48. In the event of a fault (i.e., a lost DataNode), the NameNode consults metadata, finds affected data, consults a Rack Awareness script, and instructs the DataNode to replicate. This HDFS process is described pictorially below:

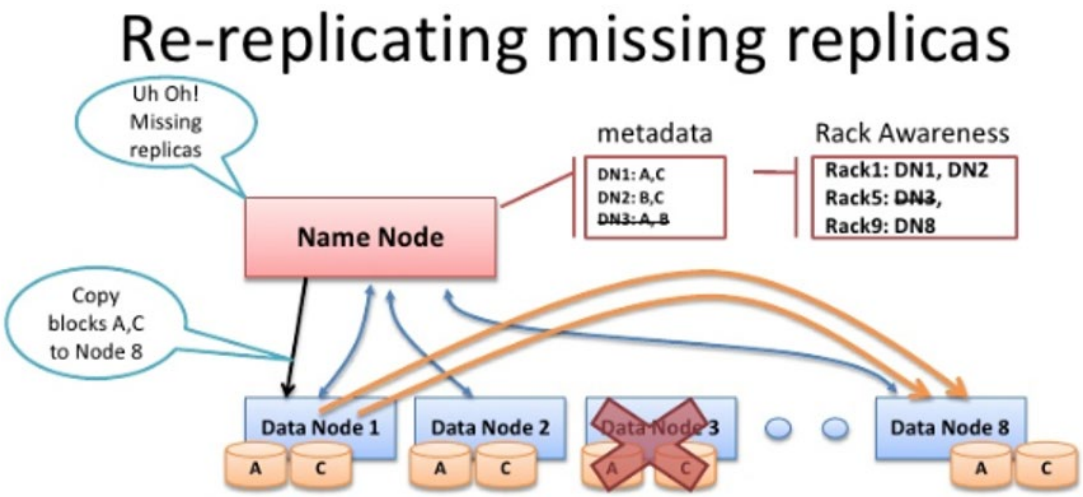

ALIBABA CLOUD

49. AliCloud has been using Hadoop since 2008.

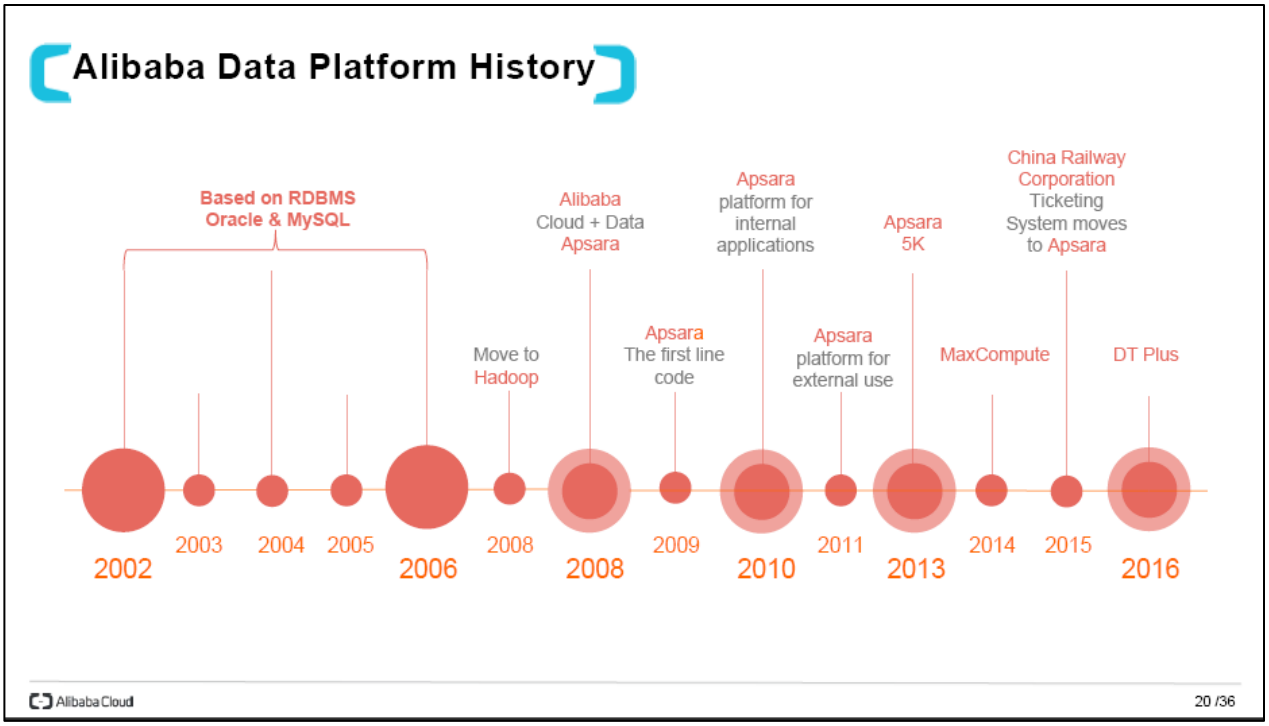

50. AliCloud describes its Hadoop Ecosystem including YARN for resource management and HDFS:

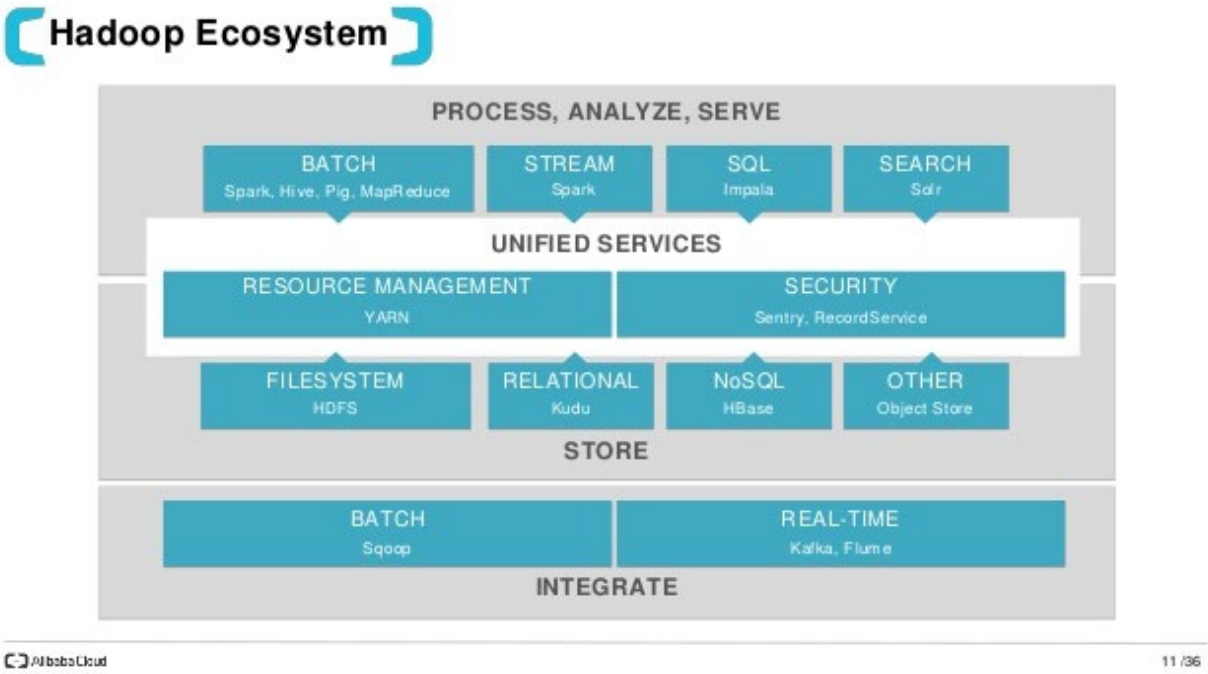

51. In an AliCloud implementation of an HDFS instance, the NameNode manages the file system namespace and regulates access to files by clients and DataNodes manage storage attached to the nodes they run on as shown below:

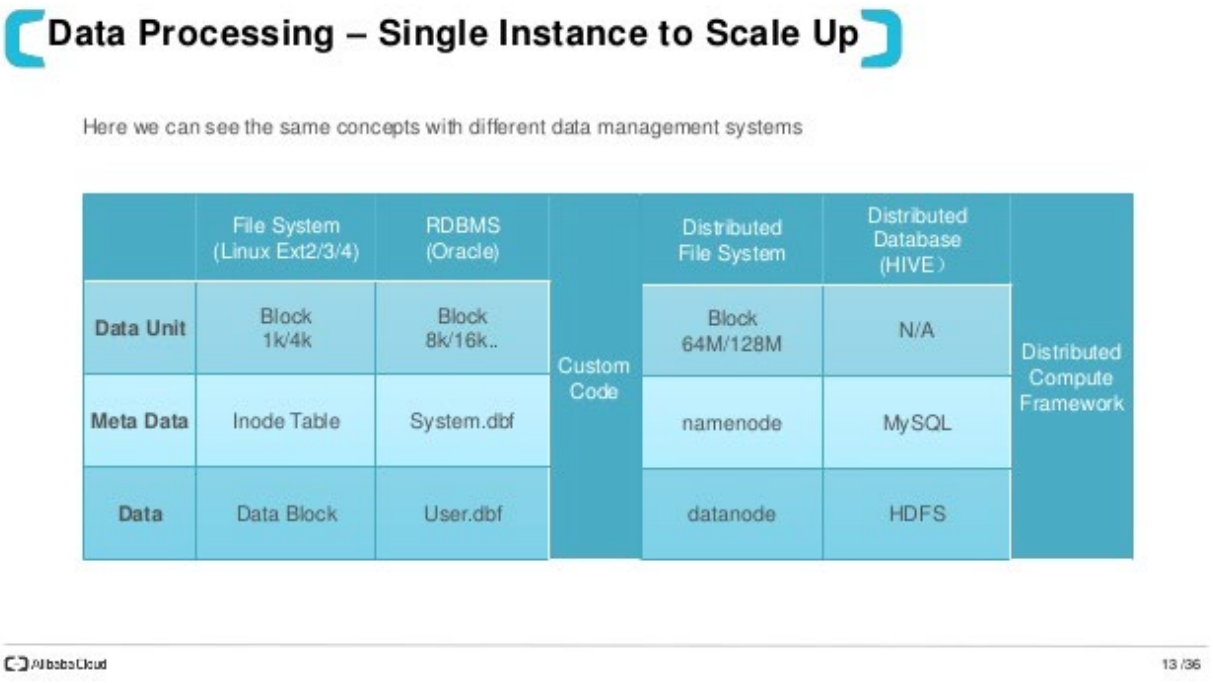

52. Since 2013, AliCloud's data infrastructure has used Hadoop YARN to serve various types of workloads such as batch jobs, streaming, machine learning, OLAP, and online services.

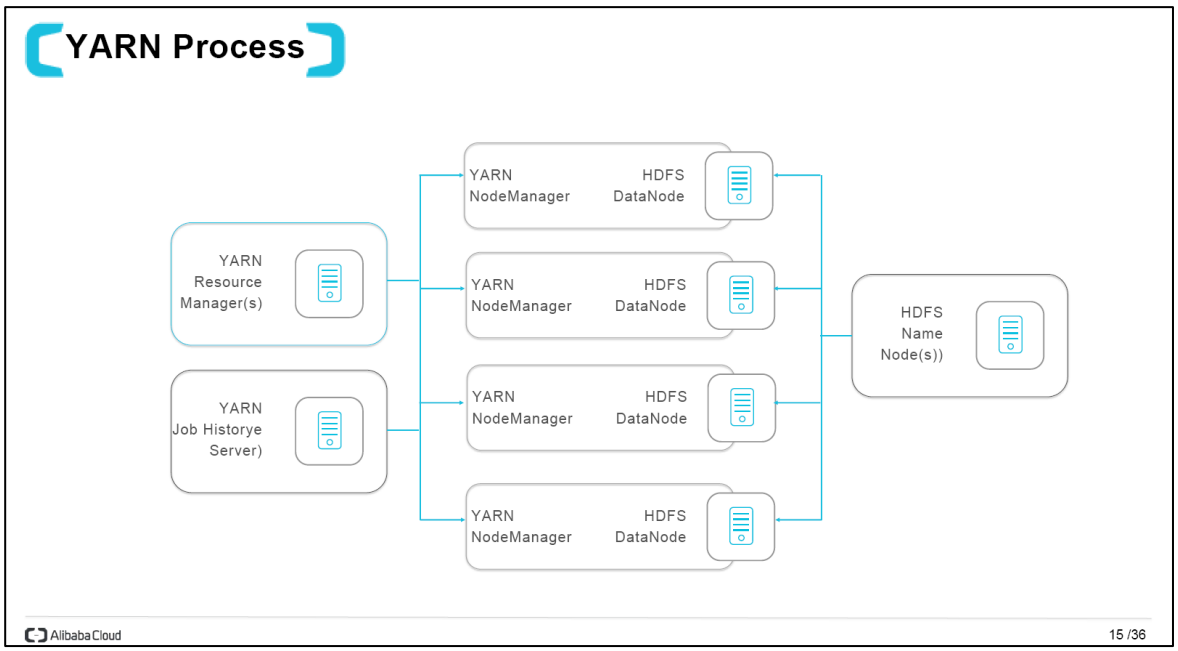

53. AliCloud's Elastic MapReduce (E-MapReduce) product is a big-data processing solution based upon Hadoop and built on the Alibaba Cloud Elastic Compute Service (ECS).

54. AliCloud describes its E-MapReduce product as "essentially the cluster service of Hadoop." E-MapReduce enables AliCloud users to provision distributed Hadoop clusters and process data in use cases such as trend analysis, data warehousing, and analysis of continuously streaming data.

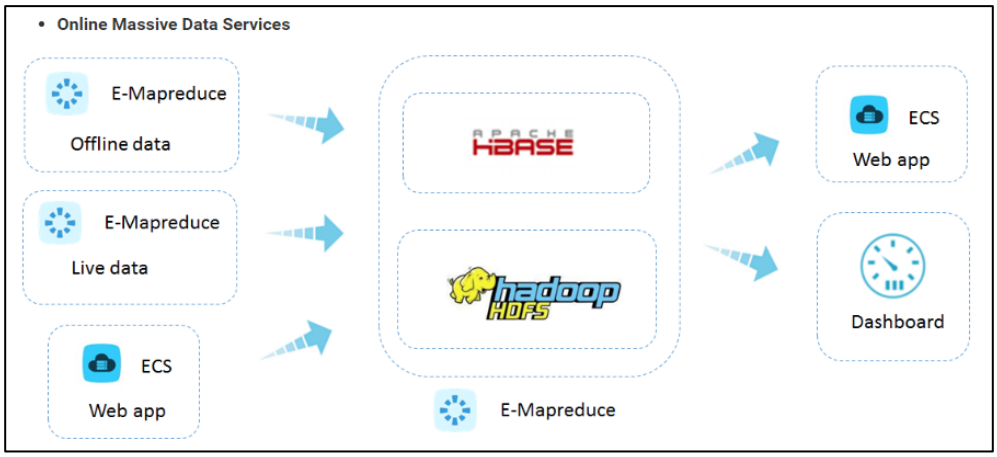

55. AliCloud E-MapReduce users can import data to and export data from AliCloud OSS (Object Storage Service) and AliCloud RDS (Relational Database Service).

56. AliCloud provides the following description of its E-MapReduce product:

Composition of E-MapReduce

The core component directly oriented to an E-MapReduce user is the cluster. An E-MapReduce cluster is a Spark and Hadoop cluster consisting of multiple ECS Alibaba Cloud instances. For example, in Hadoop, generally some daemon processes run on each ECS instance (such as namenode, datanode, resourcemanager, and nodemanager), which make up the Hadoop cluster . The nodes running namenode and resourcemanager are known as master nodes, while those running datanode and nodemanager are called slave nodes.

Alibaba Cloud E-MapReduce Product Information, *available at* https://www.alibabacloud.com/help/doc-detail/28068.htm.

57. The figure below shows an AliCloud E-MapReduce cluster consisting of one master node and three slave nodes:

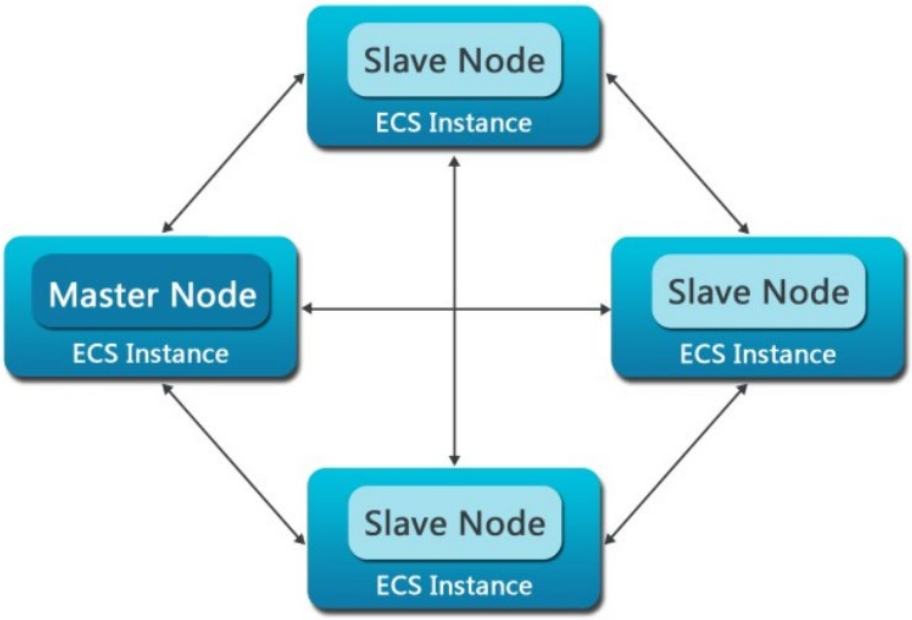

Alibaba Cloud E-MapReduce Product Information at 2.

58. An AliCloud user can apply for cluster resources according to the geographical location where the user or data source is located including at AliCloud's US East 1 and US West 1 datacenters.

59. After AliCloud subscribers select an ECS model, the entire cluster is created within a matter of minutes. Once it is online, the AliCloud Hadoop cluster dynamically adds nodes as needed.

60. AliCloud E-MapReduce operates on data stored on HDFS and supports a variety of computational processes including machine learning, process orchestration, stream processing, and graph analytics.

61. AliCloud's E-MapReduce Storage Guide describes two types of disks available for data storage: cloud disks and ephemeral disks. In E-MapReduce, when the hosting node is released, data in all cloud and ephemeral disks is cleared. Hadoop HDFS uses all data disks for data storage, and Hadoop YARN also uses all data disks as on-demand data storage for computing.

62. In AliCloud E-MapReduce, OSS can be used as HDFS. HDFS commands directly handle OSS data.

HDFS commands directly handle OSS data

```
hadoop fs -ls oss://bucket/path
hadoop fs -cp hdfs://user/path  oss://bucket/path
```

63. AliCloud provides pricing information to users to purchase E-MapReduce Hadoop clusters:

> In the created cluster, 8 Mbps public network bandwidth will be opened for the Master nodes (for HA Hadoop cluster, both of the two Master nodes will have the 8 Mbps bandwidth). The traffic will be paid as you go, which is not included in the cost of the Hadoop clusters. It only charges outflow traffic on an hourly basis, while the inflow traffic is free of charge. For example, if you use 10 GB of outbound public traffic in an hour, the charge will be 10GB * price per GB (dollar/h). The traffic fees in different regions are slightly different.

64. AliCloud publishes the following the E-MapReduce product architecture:

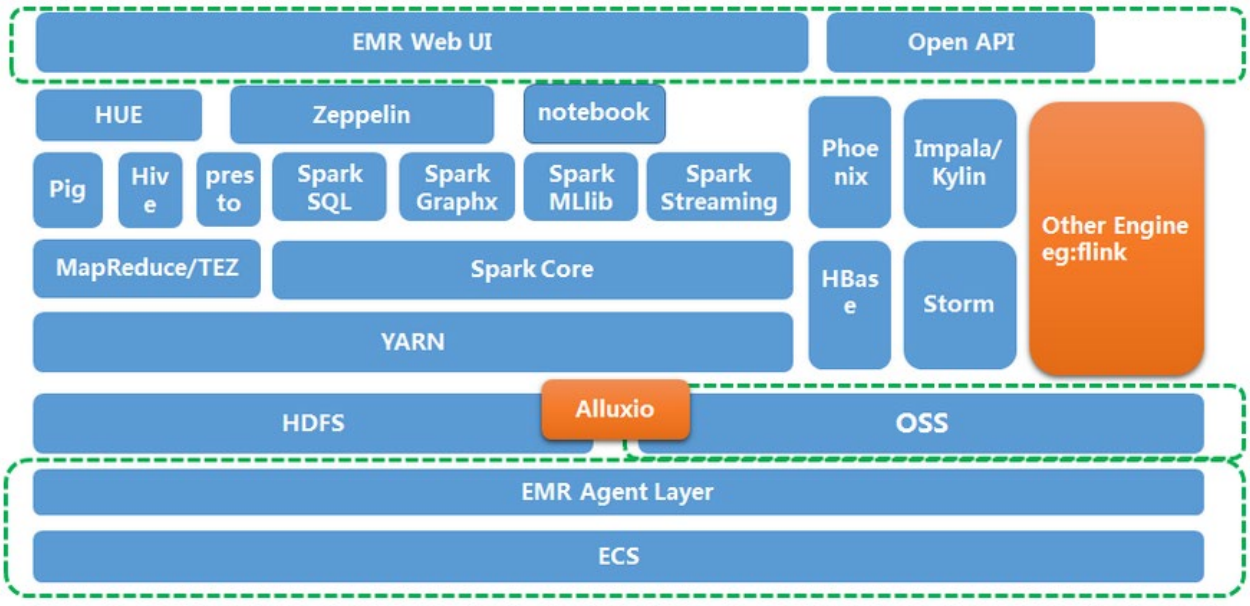

https://www.alibabacloud.com/help/doc-detail/28070.htm.

65. AliCloud instructs users how to configure and use HDFS to interface with and integrate with AliCloud products and services.

HDFS data source config

Last Updated: Apr 03, 2018

HDFS, as a distributed file system, allows both reading and writing data, and supports configuring synchronization in script mode.

Procedure

1. Log on to the **DataWorks console** as a developer, and click **Enter Project**, as shown in the following figure.

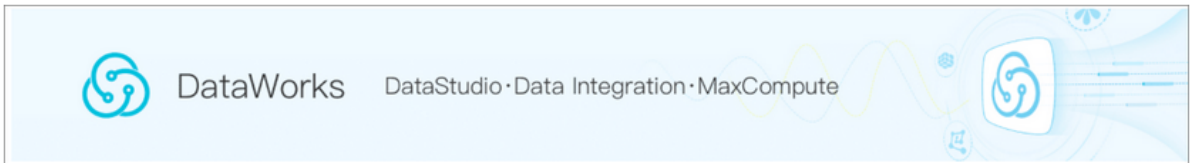

https://www.alibabacloud.com/help/doc-detail/53541.htm

Configure HDFS writer

Last Updated: Apr 03, 2018

HDFS Writer is used to write TextFile, ORCFile, and ParquetFile to the specified path in the HDFS file system. The files can be associated with Hive tables. You must configure the data source before configuring the HDFS Writer plug-in. For more information, see HDFS data source config.

How to implement HDFS Writer:

1. Create a temp directory that does not exist in the HDFS file system based on the path you specified. Naming rule: path_random 1. Write the files that have been read to this temp directory.

https://www.alibabacloud.com/help/faq-detail/53292.htm.

Configure HDFS reader

Last Updated: Apr 12, 2018

HDFS Reader provides the ability to read the data stored by distributed file systems. At the underlying implementation level, HDFS Reader retrieves the file data on the distributed file system, coverts the data into Data Integration transport protocol and transfers it to the Writer.

HDFS Reader provides the ability to read file data from the Hadoop distributed file system HDFS and coverts the data into Data Integration transport protocol.

https://www.alibabacloud.com/help/faq-detail/53291.htm.

66. AliCloud's Apsara Stack is a full-stack cloud solution created by Alibaba Cloud for medium- and large-size enterprise-class customers. AliCloud Apsara features E-MapReduce as a "Big Data Service."

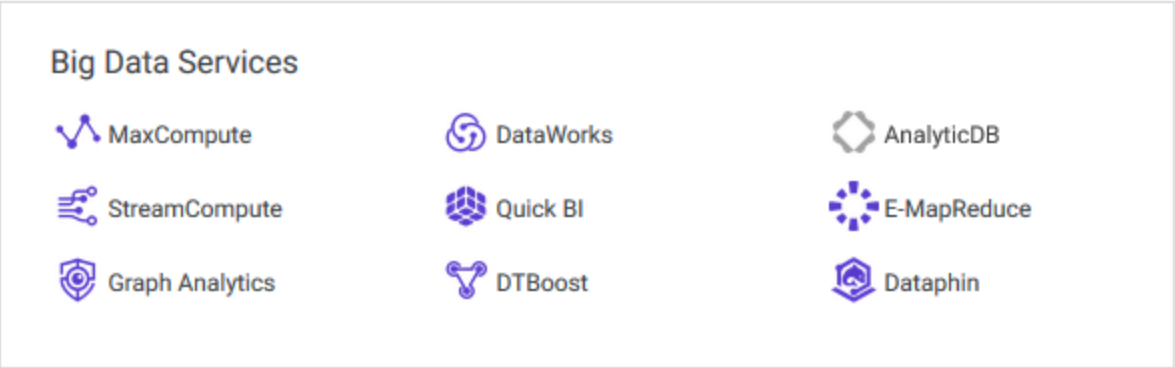

https://www.alibabacloud.com/product/apsara-stack.

67. AliCloud enables integration of its Hadoop E-MapReduce product and its MaxCompute data processing platform.

> Now, the MaxCompute platform provides a plug-in that allows you to adapt Hadoop MapReduce code to MaxCompute MapReduce specifications. MaxCompute offers a degree of flexibility regarding binary-level compatibility for Hadoop MapReduce jobs. It means that, without modifying the code, you can specify configurations to directly run original Hadoop MapReduce Jar packages on MaxCompute. Download the development plug-in to get started. This plug-in is currently in the testing stage, therefore, does not support custom comparators or key types.

https://www.alibabacloud.com/help/doc-detail/44626.htm?spm=a2c63.p38356.a1.5.37214a856KhfS6.

68. In addition to providing HDFS and YARN products and services to AliCloud customers, AliCloud uses them internally in connection with Alibaba search, advertising, security, and other Alibaba applications.

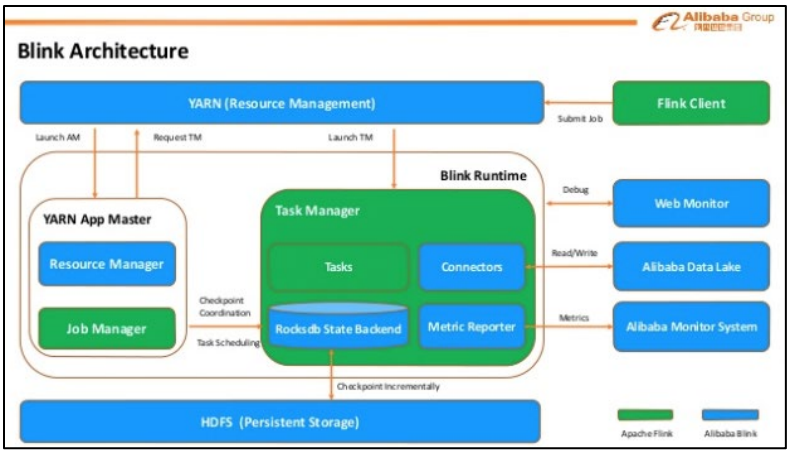

COUNT I
INFRINGEMENT OF U.S. PATENT NO. 7,543,177

69. MNS re-alleges and incorporates by reference the preceding paragraphs as if stated here.

70. The Alibaba defendants directly and indirectly infringe at least claims 1 and 13 of the ’177 Patent.

71. The Alibaba defendants make, use, sell, offer for sale, and/or import Alibaba Cloud Products and Services including those specifically identified herein and categorized by AliCloud as Elastic Computing, Storage and Content Delivery, Database Services, Analytics and Big Data, and the Apsara Stack.

72. In particular, these Accused Instrumentalities include the AliCloud HDFS storage system and YARN resource management service, E-MapReduce product and service, AliCloud Elastic Computing Service, and AliCloud Apsara featuring E-MapReduce.

73. The Accused Instrumentalities embody and practice the subject matter claimed in the asserted claims of the ’177 Patent.

74. Asserted claim 1 of the ’177 Patent recites a storage system, comprising: one or more memory sections, including: one or more memory devices having storage locations for storing data, and a memory section controller capable of detecting faults in the memory section and transmitting a fault message in response to the detected faults; one or more switches, including: one or more interfaces for connecting to one or more external devices; a switch controller that executes software, including a routing algorithm; and a selectively configurable switch fabric connected to one or more memory sections and the one or more interfaces and interconnecting the memory sections and the one or more interfaces based on the routing algorithm stored in the switch controller; and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm, and wherein the management system is further capable of determining and

changing the routing algorithm for use by the selectively configurable switch fabric in interconnecting the memory sections and the one or more interfaces, providing the determined routing algorithm to the switch controller, and instructing the switch controller to execute the determined routing algorithm.

75. Asserted claim 13 of the '177 Patent recites a method for use in a storage system, comprising: storing data in a storage locations in a memory device, the memory device included in a memory section; a management system determining a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface; the management system providing the determined routing algorithm to the switch controller and instructing the switch controller to execute the determined routing algorithm; the selectively configurable switch connecting the memory section to the interface based on the routing algorithm; detecting by a memory section controller a fault in regard to the data stored in the memory device and transmitting a fault message in response to the detected fault to the management system; receiving the fault message at the management system; and the management system removing from service the memory section from which the fault message was received by changing the routing algorithm.

76. The Accused Instrumentalities, and HDFS implementations on AliCloud, are storage systems.

77. A typical architecture of a Hadoop cluster features Slave nodes for storage and the NameNode that oversees and coordinates the data storage function.

78. In normal operation, the Accused Instrumentalities implementing HDFS store data blocks in a DataNode's (memory section) local file system that uses storage including memory devices (e.g., HDD, SSD). The memory devices store data in physical storage locations (e.g., HDD sectors, SSD blocks).

79. The Accused Instrumentalities include a management system that determines a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory

section and an interface.

80. In normal operation, the Accused Instrumentalities implementing HDFS manage the HDFS NameSpace (e.g., by operation of the HDFS NameNode daemon) and map data file names to sets of data blocks, map data blocks to specific DataNodes, and map DataNodes to specific racks in the HDFS cluster.

81. In the Accused Instrumentalities, NameNode NameSpace tables and resultant NameNode instructions based on them (i.e. the I/O path a HDFS client uses to read/write a specific data block) are routing algorithms used by the HDFS NameNode (switch controller) that controls how specific HDFS I/O requests traverse the HDFS cluster.

82. Consistent with the asserted claims, the Accused Instrumentalities implementing HDFS achieve high fault tolerance by ensuring persistence of file system metadata.

83. In the accused HDFS instances, the HDFS namespace is stored by the NameNode, which uses a transaction log called the EditLog to persistently record every change that occurs to file system metadata.

84. For example, creating a new file in an AliCloud HDFS instance causes the NameNode to insert a record into the EditLog. Changing the replication factor of a file also causes a new record to be inserted into the EditLog. The NameNode stores the EditLog, and the entire file system NameSpace, including the mappings and system properties, is stored by the NameNode.

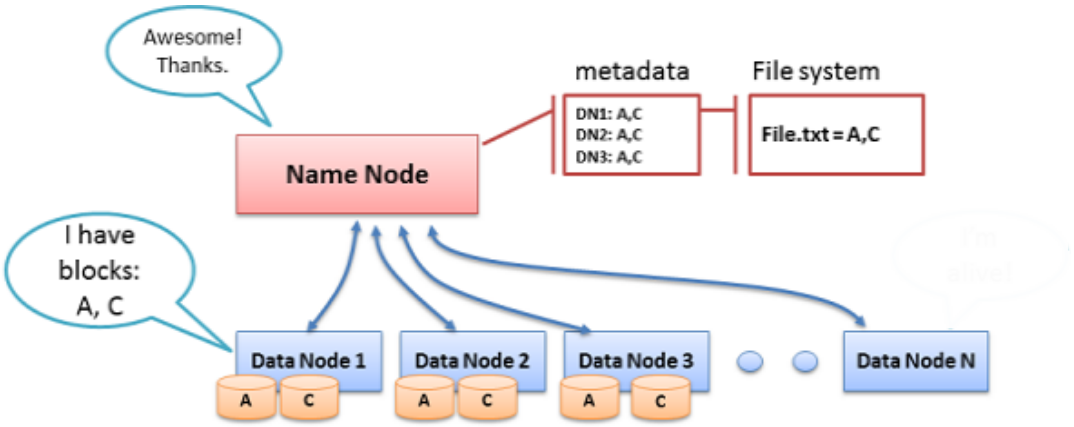

85. The figure above provides a representative diagram of the switch controller

(NameNode), routing algorithm (metadata and file system), memory section (DataNode N), and memory devices (devices labeled A, C).

86. In the accused HDFS implementations, the NameNode daemon determines the routing algorithm by processing the metadata tables in response to HDFS client Read and Write operations (exemplified in the figures below).

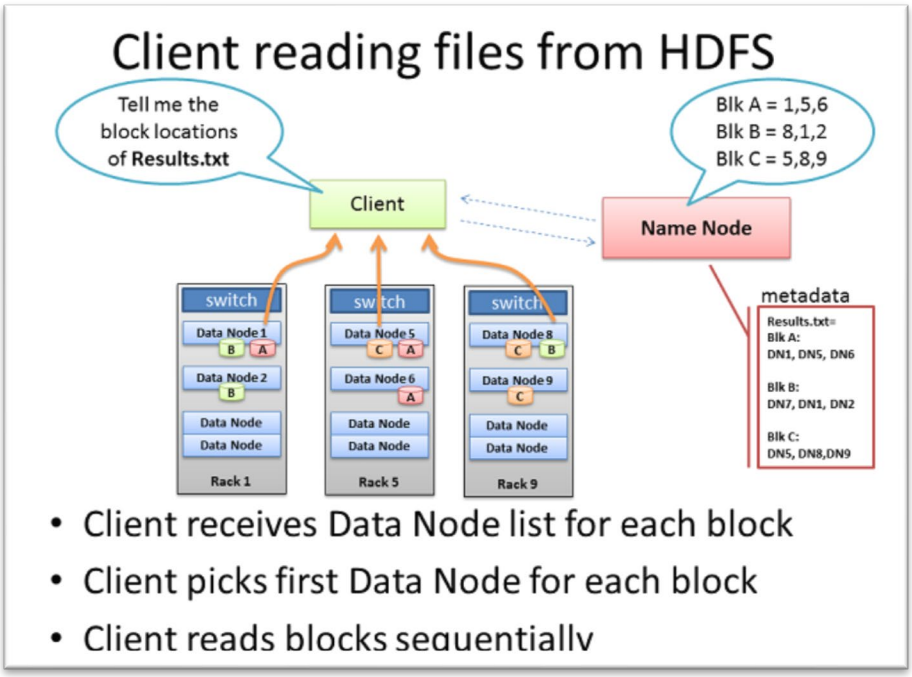

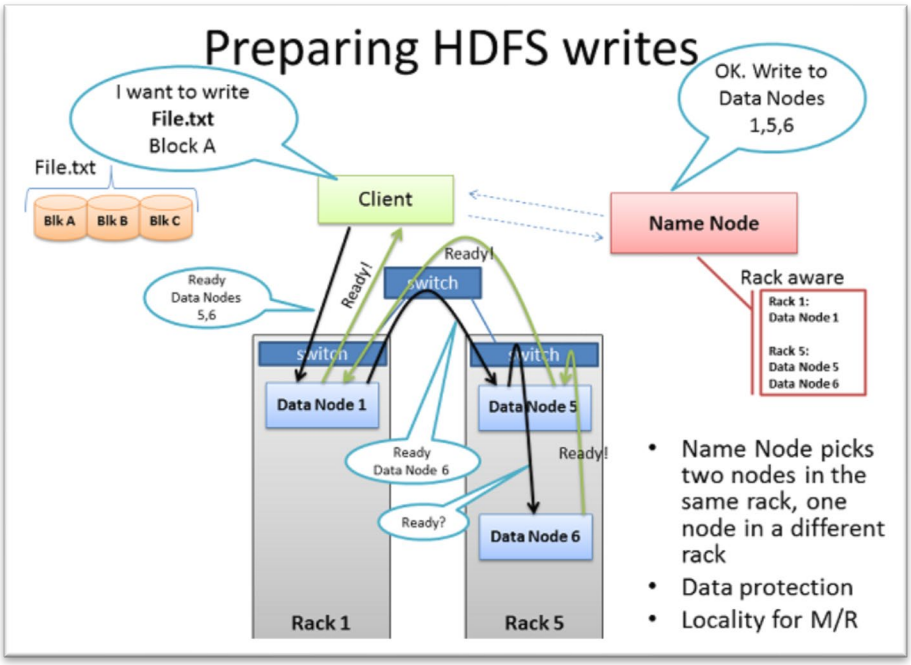

87. The Accused Instrumentalities include network switches.

88. In large clusters, the Accused Instrumentalities spread the nodes across multiple racks. Nodes of a rack share a switch, and these rack switches, which are selectively configurable, are in turn connected by one or more core switches.

89. In the Accused Instrumentalities, selectively configurable rack switches connect HDFS data nodes (memory sections) to an interface.

90. In the event of an HDFS I/O request, the rack switch routes the request to the proper HDFS data node in accordance with the HDFS file system NameSpace that includes the mapping of blocks to files.

91. In normal operation of the Accused Instrumentalities, a memory section controller (e.g., data node daemon) detects a fault in regard to data stored in the memory device and a fault message is transmitted to the management system (e.g., HDFS NameNode) in response to the detected fault.

92. During normal operation, each DataNode periodically sends a heartbeat message to the NameNode. If a subset of DataNodes lose connectivity with the NameNode, the NameNode detects the fault by the absence of a heartbeat message and marks the affected DataNodes as dead and ceases forwarding any new I/O requests to them. The NameNode tracks which blocks need to be replicated due to a fault and initiates replication when necessary.

93. By default, the heartbeat is transmitted every three seconds, set by dfs.heartbeat.interval.

94. In addition to detecting a fault by monitoring heartbeats, HDFS DataNodes create threads that run a DataBlockScanner object that scans the data blocks (and replicas) stored in the DataNode to detect faults.

95. The Name Node daemon receives the fault message in the NameNode (management system) due to either a disruption in heartbeats from a DataNode or receipt of a DataBlockScanner report indicating a fault.

96. During normal operation of the Accused Instrumentalities, upon detecting a dead DataNode (e.g., a DataNode with no heartbeat) the NameNode daemon (management

system) bypasses the dead DataNode and instead sends I/O requests to the other DataNodes storing replicas of blocks that were stored on the dead DataNode. If a corrupted block is detected (e.g., via DataBlockScanner) the NameNode daemon (management system) marks the block replica as corrupt and then schedules a copy of the block to be replicated on another DataNode, which results in an updated HDFS NameSpace (a new routing algorithm) so its replication factor is back at the expected level. Thus, during normal operation, the management system removes from service the memory section from which a fault message was received by changing the routing algorithm.

97. Through technical support and publication of instructional information, Defendants encourage, aid, and direct end users of the Accused Instrumentalities to use and operate them, consistent with AliCloud's instructions, to perform the asserted method claims.

98. Defendants are on notice of the infringing products, services, features, and how end-users of the Accused Instrumentalities operate them to perform the claimed methods and use the claimed apparatuses.

99. Defendants' infringing conduct has damaged MNS.

100. Defendants are liable to MNS in an amount that adequately compensates it for Defendants' infringement, which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 7,958,388

101. MNS re-alleges and incorporates by reference the preceding paragraphs as if stated here.

102. The Alibaba defendants directly and indirectly infringe at least claims 1 and 2 of the '388 Patent.

103. The Alibaba defendants make, use, sell, offer for sale, and/or import Alibaba Cloud Products and Services including those specifically identified herein and categorized by AliCloud as Elastic Computing, Storage and Content Delivery, Database Services,

Analytics and Big Data, and the Apsara Stack.

104. In particular, these Accused Instrumentalities include the AliCloud HDFS storage system and YARN resource management service, E-MapReduce product and service, AliCloud Elastic Computing Service, and AliCloud Apsara featuring E-MapReduce.

105. The Accused Instrumentalities embody and practice the asserted claims of the '388 Patent.

106. Asserted claim 1 of the '388 Patent recites a storage system, comprising: one or more memory sections, including one or more memory devices having storage locations for storing data, and a memory section controller capable of detecting faults in the memory section and transmitting a fault message in response to the detected faults; one or more switches, including one or more interfaces for connecting to one or more external devices; a switch controller that executes software, including a routing algorithm; and a selectively configurable switch fabric connected to one or more memory sections and the one or more interfaces and interconnecting the memory sections and the one or more interfaces based on the routing algorithm; and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm, and wherein the management system is further capable of determining the routing algorithm for use by the selectively configurable switch fabric in interconnecting the memory sections and the one or more interfaces, and providing the routing algorithm to the switch controller.

107. Asserted claim 2 of the '388 Patent recites a method for use in a storage system, comprising: storing data in storage locations in a memory device, the memory device included in a memory section; determining, by a management system, a routing algorithm for use by a switch controller that executes software, including the routing algorithm; providing, by the management system, the routing algorithm to the switch controller; executing, by the switch controller, the routing algorithm, to configure a configurable switch connecting the memory section to an interface; detecting a fault

associated with the data in the storage locations in the memory device; determining, by the management system in response to the detecting, a new routing algorithm that redirects data for the memory device to a replacement memory device; and providing the new routing algorithm to the switch controller.

108. In normal operation of the Accused Instrumentalities, the management system determines a new routing algorithm that redirects data for the memory device to a replacement memory device in response to detecting a fault.

109. During normal operation and upon detecting a dead DataNode (e.g., a DataNode with no heartbeat) the NameNode daemon (management system) bypasses the dead DataNode and sends I/O requests to other DataNodes storing replicas of blocks that were stored on the dead DataNode. The NameNode then schedules creation of new block replicas (to be stored on replacement memory devices) which result in an updated HDFS NameSpace (new routing algorithm).

110. Upon detecting a corrupted block (via DataBlockScanner) the NameNode daemon (management system) marks the block replica as corrupt and then schedules a copy of the block to be replicated (stored on replacement memory devices) on another datanode, so its replication factor is back at the expected level. This results in an updated HDFS NameSpace (new routing algorithm).

111. During normal operation, the DataBlockScanner object creates a list of replicas that serves as the initial list of data blocks that it will scan for errors. When the NameNode becomes aware that a block is corrupt, it updates its internal tables to indicate that a block on a specific DataNode is corrupt and enters the corrupt replica into a list of blocks needing additional replicas. Once the replica has been created, the identity of the new replicas in this DataNode are sent to the NameNode.

112. When the NameNode daemon detects a fault (e.g. a dead NameNode or corrupt data block) an updating of the HDFS NameSpace is triggered that results in updates to the NameNode NameSpace (a new routing algorithm provided to the switch controller).

113. Through technical support and publication of instructional information,

Defendants encourage, aid, and direct end users of the Accused Instrumentalities to use and operate them, consistent with AliCloud's instructions, to perform the asserted method claims.

114. Defendants are on notice of the infringing products, services, features, and how end-users of the Accused Instrumentalities operate them to perform the claimed methods and use the claimed apparatuses.

115. Defendants' infringing conduct has damaged MNS.

116. Defendants are liable to MNS in an amount that adequately compensates it for Defendants' infringement, which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

PRAYER FOR RELIEF

MNS prays for the following relief:

a) A judgment be entered that Defendants have directly and indirectly infringed one or more claims of the Asserted Patents;

b) A judgment be entered that the Asserted Patents are valid and enforceable;

c) An award of damages adequate to compensate MNS for Defendants' infringement up until the date such judgment is entered, including prejudgment and post-judgment interest, costs, and disbursements as justified under 35 U.S.C. § 284 and an accounting, if necessary to adequately compensate MNS for Defendants' infringement;

d) A judgment that MNS be awarded attorneys' fees, costs, and expenses incurred in prosecuting this action; and

e) A judgment that MNS be awarded such further relief at law or in equity as the Court deems just and proper.

DEMAND FOR JURY TRIAL

MNS demands trial by jury for all issues so triable pursuant to Fed. R. Civ. P. 38(b) and Civil L.R. 3-6(a).

Dated: January 15, 2019

By /s/ *Stephen M. Lobbin*

Cabrach J. Connor
Texas Bar No. 24036390 (*pro hac vice* pending)
cab@connorkudlaclee.com
Kevin S. Kudlac
Texas Bar No. 00790089 (*pro hac vice* pending)
kevin@connorkudlaclee.com
Jennifer Tatum Lee
Texas Bar No. 24046950 (*pro hac vice* pending)
jennifer@connorkudlaclee.com
CONNOR KUDLAC LEE PLLC
609 Castle Ridge Road, Suite 450
Austin, TX 78746
Telephone: 512.777.1254
Facsimile: 888.387.1134

Stephen M. Lobbin
California Bar No. 181195
sml@smlavvocati.com
SML AVVOCATI P.C.
888 Prospect Street, Suite 200
San Diego, CA 92037
Telephone: 949.636.1391

Attorneys for Plaintiff
Mobile Networking Solutions, LLC