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* forthcoming)
cab@connorkudlaclee.com
KEVIN S. KUDLAC
Texas Bar No. 00790089 (*pro hac vice* forthcoming)
kevin@connorkudlaclee.com
JENNIFER TATUM LEE
Texas Bar No. 24046950 (*pro hac vice* forthcoming)
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 NETWORK SOLUTIONS, LLC

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

MOBILE NETWORK SOLUTIONS, LLC, Plaintiff, vs. UBER TECHNOLOGIES, INC., Defendant.	CASE NO. 3:19-cv-219_______ **COMPLAINT FOR PATENT INFRINGEMENT** **DEMAND FOR JURY TRIAL**

Mobile Network Solutions, LLC ("MNS") files this Original Complaint for Patent Infringement against Uber Technologies, Inc. ("Uber") 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. Uber is a corporation incorporated under the laws of the State of Delaware.

3. Uber is headquartered and has a principal place of business at 1455 Market Street, 4th Floor, San Francisco, California 94103 and may be served through its agent for service of process, CT Corporation System.

4. Uber is a transportation network company known for its ride-sharing application and expansion into trucking (Uber Freight), food delivery (Uber Eats), and healthcare transportation (Uber Health). Uber relies heavily on data analytics to forecast rider demand and address bottlenecks in its operations. As Uber's business has grown, its big-data ecosystem has expanded commensurately. Uber claims to have a dataset exceeding 100 petabytes that is available to hundreds of users including Uber operations teams, data scientists and analysts, and engineering teams.

5. Several data sources feed into Uber's data ecosystem:

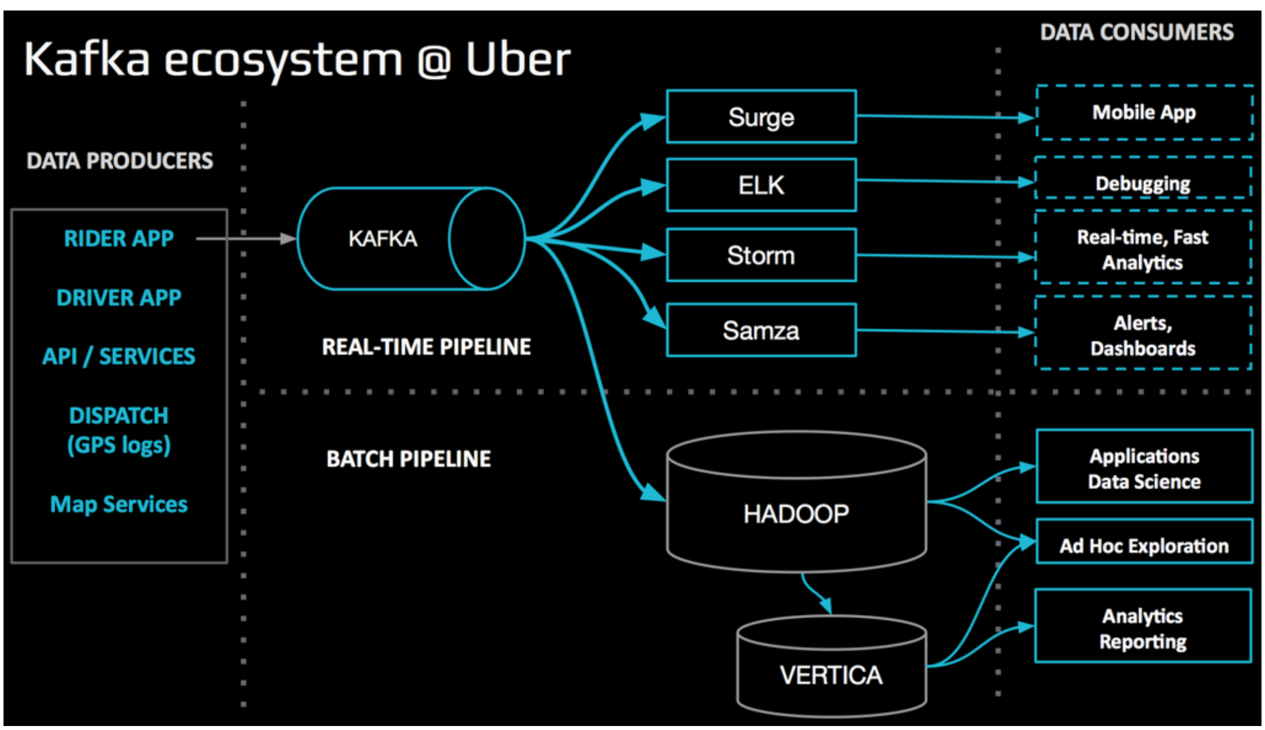

6. Uber values low latency and reliable access to its data.

7. Uber relies on reliable real-time analytics in many of its big-data use cases:

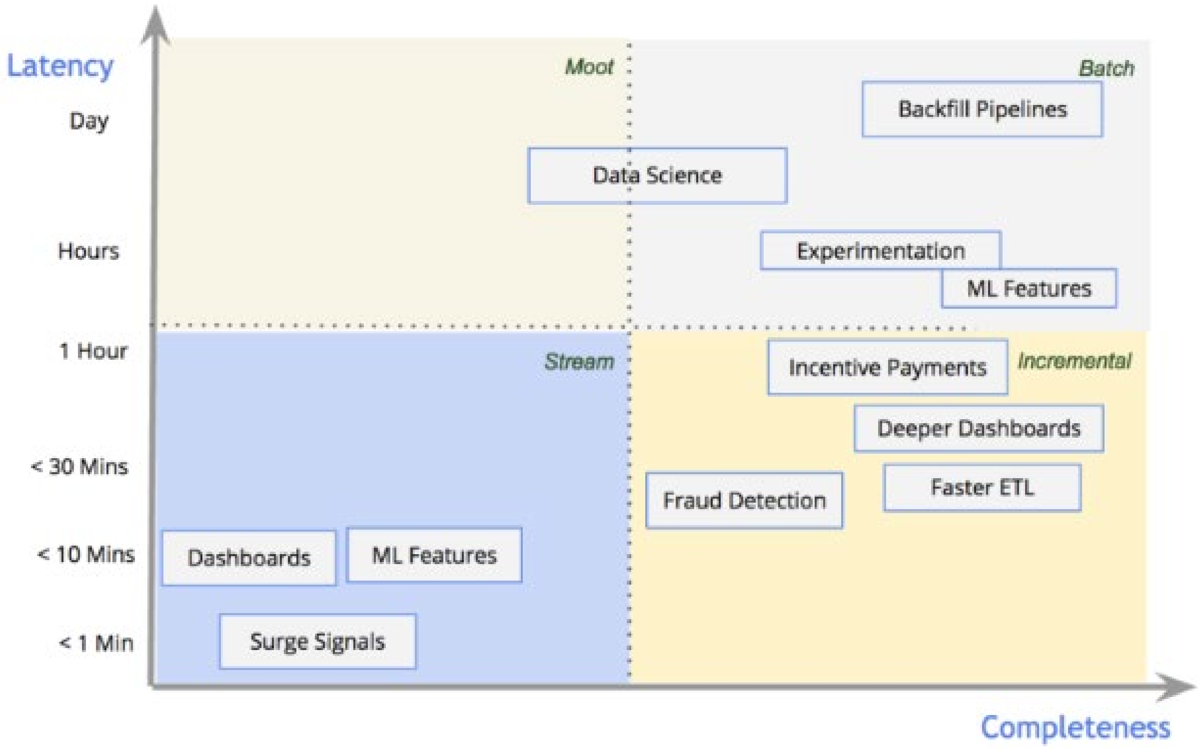

Figure 4: The above diagram demonstrates the distribution of use-cases across different latencies and completeness levels at Uber.

8. Upon information and belief, Uber operates multiple data centers in the United States including in this district, Texas, Virginia, and Colorado.

9. Uber's U.S. data centers serve domestic users (including internal Uber analysts and engineers) and store massive quantities of data generated from serving customers and users of the Uber applications in the United States.

JURISDICTION AND VENUE

10. 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 Uber 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" or the "MNS patents").

11. This Court has subject-matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and

1338(a).

12. Uber is subject to general and specific personal jurisdiction of this Court based upon its regular conduct of business in California and in this judicial district giving rise to this action and maintenance of a principal place of business in this district.

13. Uber directly and through subsidiary business units has committed and continues 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.

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

15. Venue is proper in this district pursuant to 28 U.S.C. § 1400(b) because Uber maintains a regular and established place of business in this judicial district and has committed acts of infringement in this district.

THE MNS PATENTS

16. 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."

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

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

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

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

21. 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.

22. 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.

23. 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.

24. 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.”

25. 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."

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

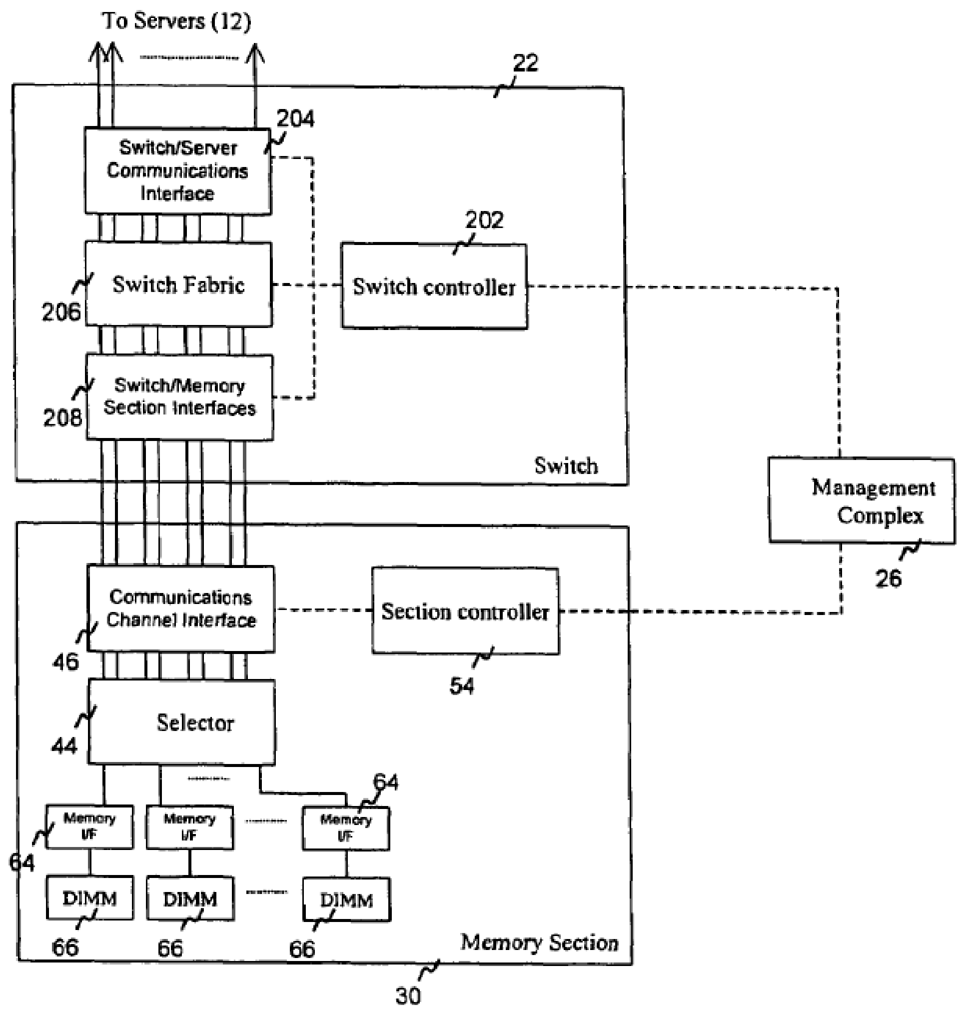

27. The inventors recognized and noted in the specification that large-scale storage systems suffered from problems in throughput for high-volume, real-time applications.

28. 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.

29. 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.

30. 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

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

32. Uber adopted Hadoop as the storage (HDFS) and compute (YARN) infrastructure for its Big Data Platform in or around 2015.

33. According to Uber, "by early 2017, [Uber's] Big Data Platform was used by engineering and operations teams across the company enabling them to access new and historical data all in one place . . . With over 100 petabytes of data in HDFS, 100,000 vcores in [Uber's] compute cluster." https://eng.uber.com/uner-big-data-platform/ (last accessed January 9, 2019).

34. Uber built a "Hadoop Data Lake" storage system to solve data infrastructure problems including those related to existing storage limitations and ETL (Extract, Transform, Load).

35. Hadoop/HDFS allows Uber to EL from data sources rather than ETL (ingesting data from multiple stores without transforming the data).

36. Since 2015, Hadoop/HDFS has been a key part of Uber's Big Data Platform infrastructure.

37. According to Uber, in 2015 it adopted Hadoop and the storage (HDFS) and compute (YARN) infrastructure for big data analysis.

38. According to Uber, "With Uber's business continuing to scale at light speed [2015-2016], we soon had tens of petabytes of data. On a daily basis, there were tens of terabytes of new data added to our data lake, and our Big Data platform grew to over 10,000 vcores with over 100,000 running batch jobs on any given day. This resulted in our Hadoop data lake becoming the centralized source-of-truth for all analytical Uber data." *Id.*

39. Uber describes the importance of its Big Data Platform in the following publication:

> *Uber's digital platform collects an incredible amount of data:* Mapping information, our movements, preferences, connections are just a few of the elements in Uber's data stores. This amount is massive, unique to Uber, and when combined creatively with other sources of data becomes a competitive weapon. Uber, like many others, is actively investing in developing additional capabilities, many of them digital; data is the critical piece underlying that strategy.

https://datacenterfrontier.com/uber-data-center-expansion/ (last accessed January 9, 2019).

40. Uber published the following diagram describing use cases for its Hadoop/HDFS Big Data Platform:

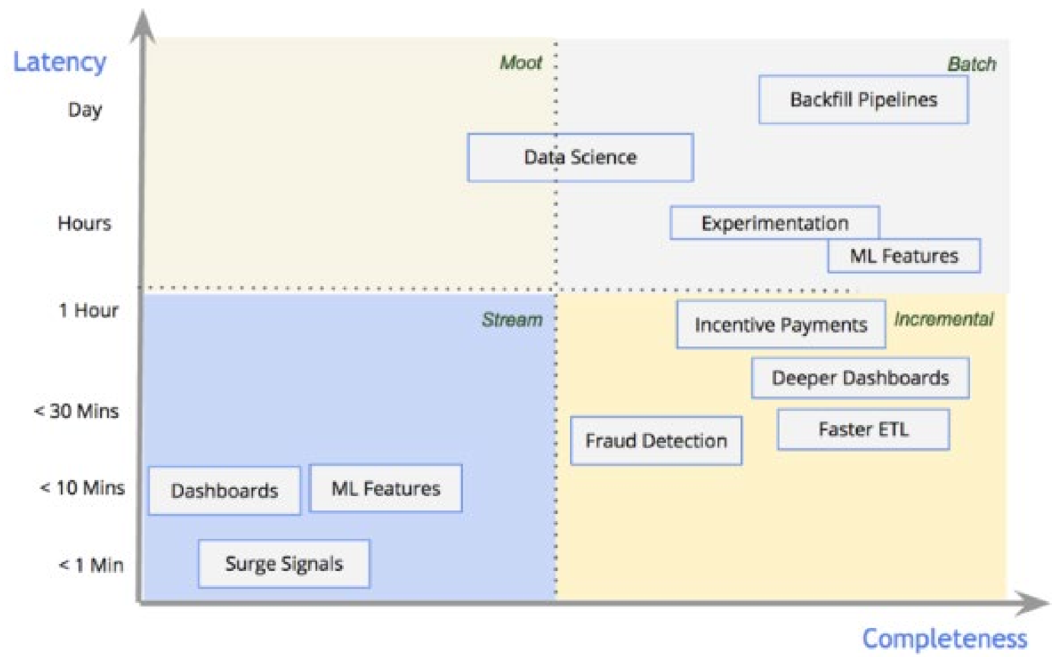

Figure 4: The above diagram demonstrates the distribution of use-cases across different latencies and completeness levels at Uber.

https://eng.uber.com/hoodie/ (last accessed January 9, 2019).

41. Uber uses Hadoop/HDFS for data analytics including geospatial computation, forecasting rider demand, fraud detection, and to address bottlenecks in the rider-partner sign-up process.

42. According to Uber, Hadoop "enables the delivery of more seamless and reliable user experiences." https://eng.uber.com/scaling-hdfs/.

43. Uber publishes the following diagram illustrating the Hadoop/HDFS architecture of its Big Data Platform:

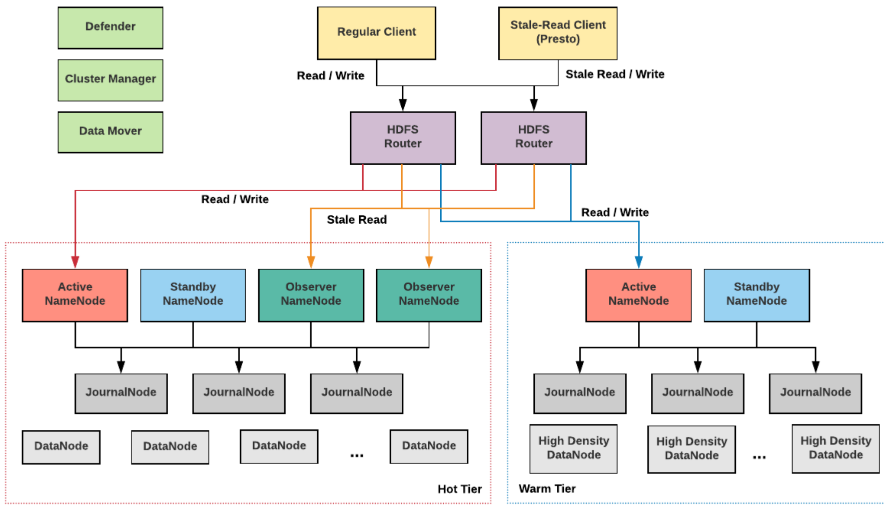

Id.

44. An HDFS deployment (e.g. an HDFS data lake) may consist of hundreds or even thousands of servers (DataNodes) that each store part of a large data file.

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

46. To handle ingestion and storage of data in Uber's Big Data Platform, Uber utilizes its HUDI (Hadoop Upsert Delete and Incremental) file format and the Marmaray framework allowing users to incrementally pull out only changed data for operations.

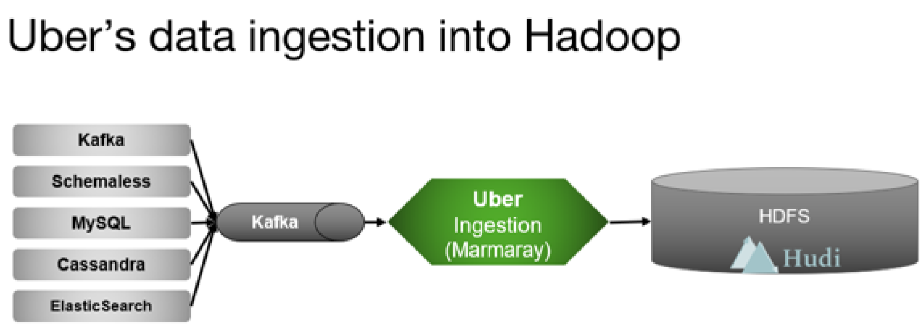

https://slideplayer .com/slide/14875299/ (last accessed January 9, 2019).

47. The graph below shows that Uber's adoption of Hadoop/HDFS resulted in a substantial decrease in pager-duty alerts even as Uber's Big Data Platform grew.

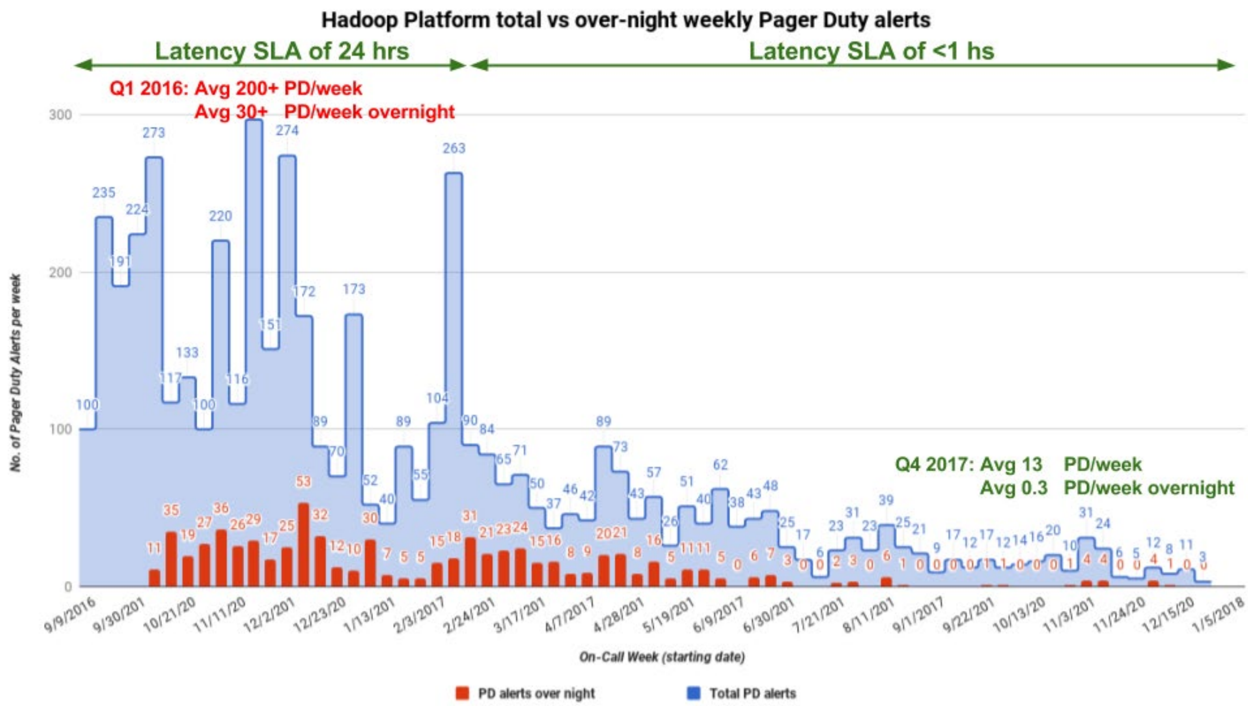

https://eng.uber.com/marmaray-hadoop-ingestion-open-source/ (last accessed January 9, 2019).

48. 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 through.

49. 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.

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

51. 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.

52. 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.

53. The NameNode is the directory tree of all files in the system and tracks where data files are kept.

54. Within the Uber Big Data Platform, Marketplace funnels real-world, real-time requests and locations into the "chutes and ladders" of Uber. Marketplace houses much of the logic for Uber products UberRUSH and UberEATS. According to Uber, Marketplace has the highest availability requirements at Uber. https://eng.uber.com/tech-stack-part-two/ (last accessed January 9, 2019).

55. Because the Marketplace stack must receive and execute in real-time, Uber states, "[e]ven brief interruptions in this area have major consequences for our users, and our business." *Id.*

56. Uber publishes the following data-flow diagram:

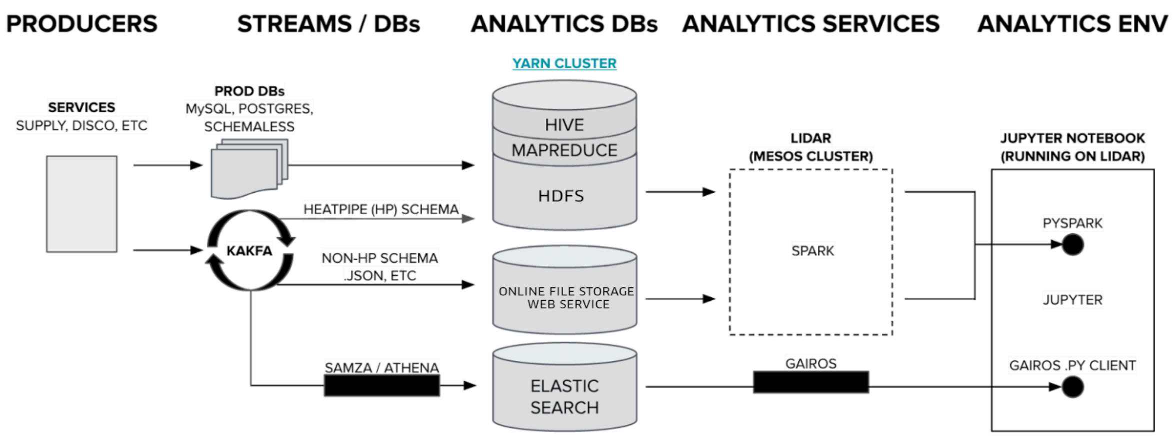

Id.

57. Uber benefits from the fault tolerance capabilities in its Hadoop/HDFS Big Data Platform, describing it as "highly available, self-healing, persistent." *Id.*

58. 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:

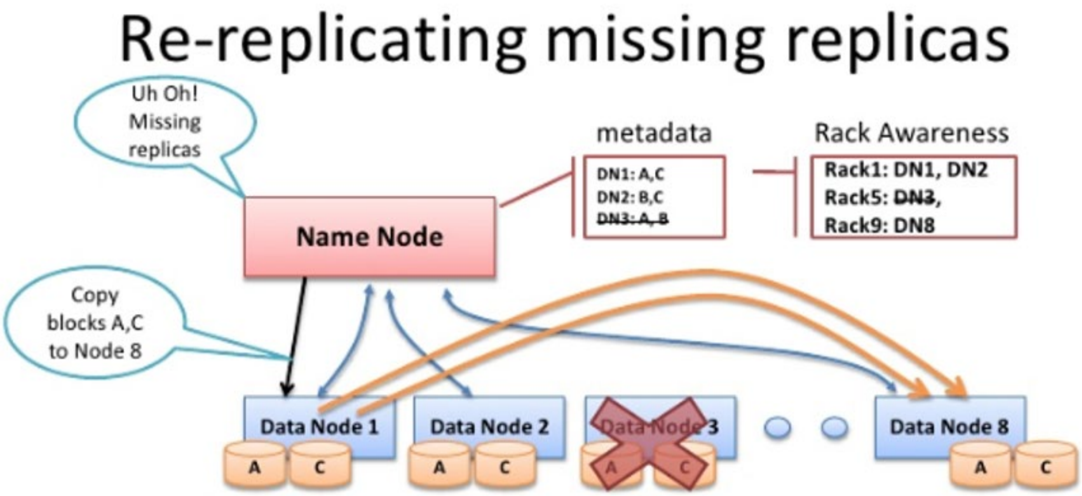

59. In Uber's implementation of HDFS, the NameNode manages the file system namespace and regulates access to files by clients and DataNodes manage storage attached to the nodes they run on.

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

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

61. Uber has and continues to infringe at least claims 1 and 13 of the '177 Patent.

62. Uber makes, uses, sells, and/or offers for sale the Uber Big-Data Platform using HDFS (the "Accused Instrumentalities").

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

64. 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.

65. 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.

66. The Accused Instrumentalities, and HDFS implementations in the Uber Big Data Platform, are storage systems.

67. A typical architecture of a Hadoop cluster features Slave nodes for storage

and the NameNode that oversees and coordinates the data storage function.

68. 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).

69. 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.

70. 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.

71. 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.

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

73. In the Accused Instrumentalities, 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.

74. For example, creating a new HDFS file in the Accused Instrumentalities 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.

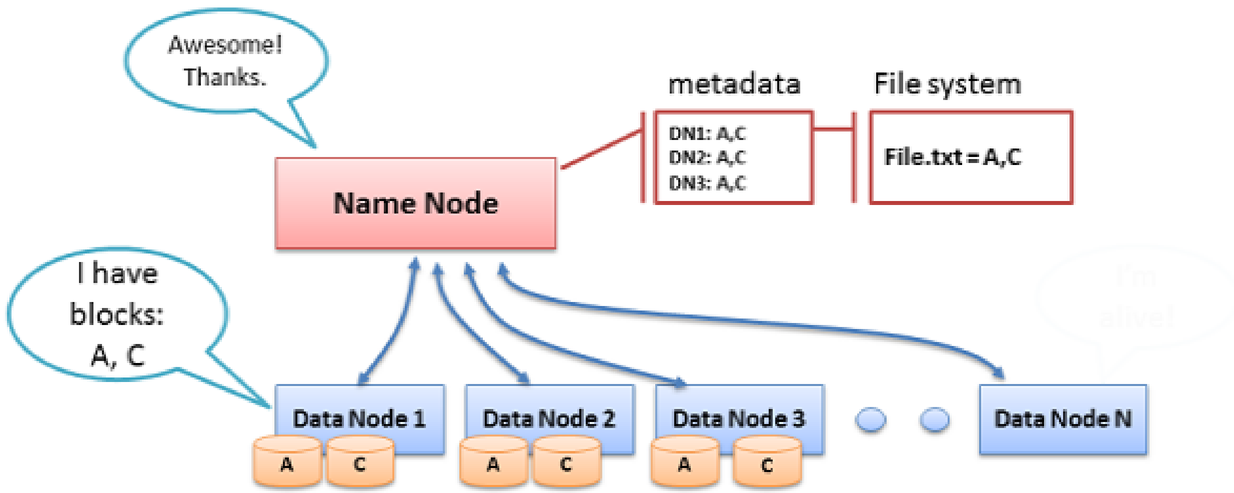

75. The figure above provides a representative diagram of the switch controller (NameNode), routing algorithm (metadata and file system), memory section (DataNode N), and memory devices (devices labeled A, C).

76. 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).

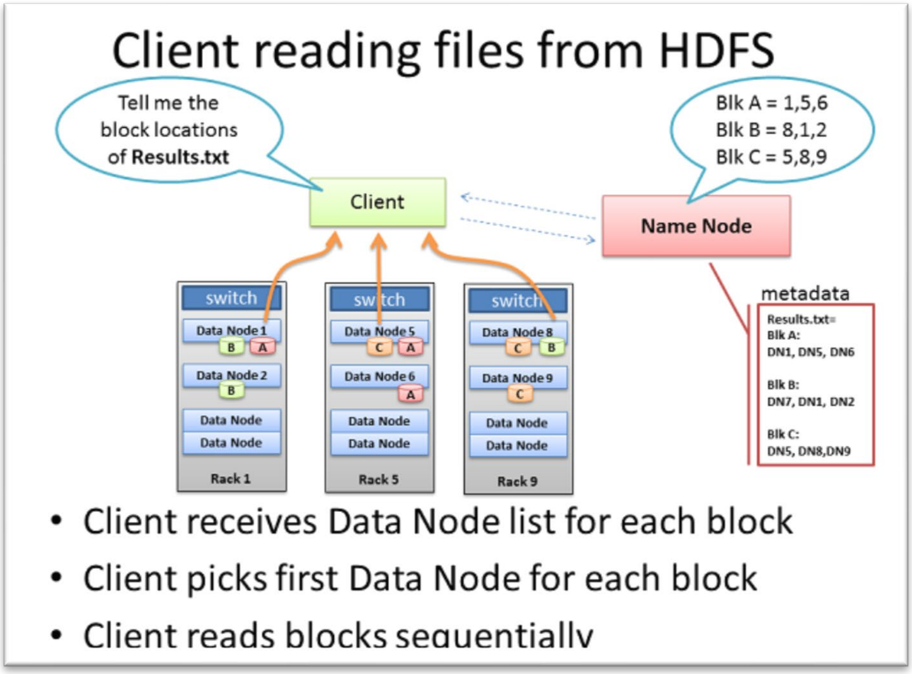

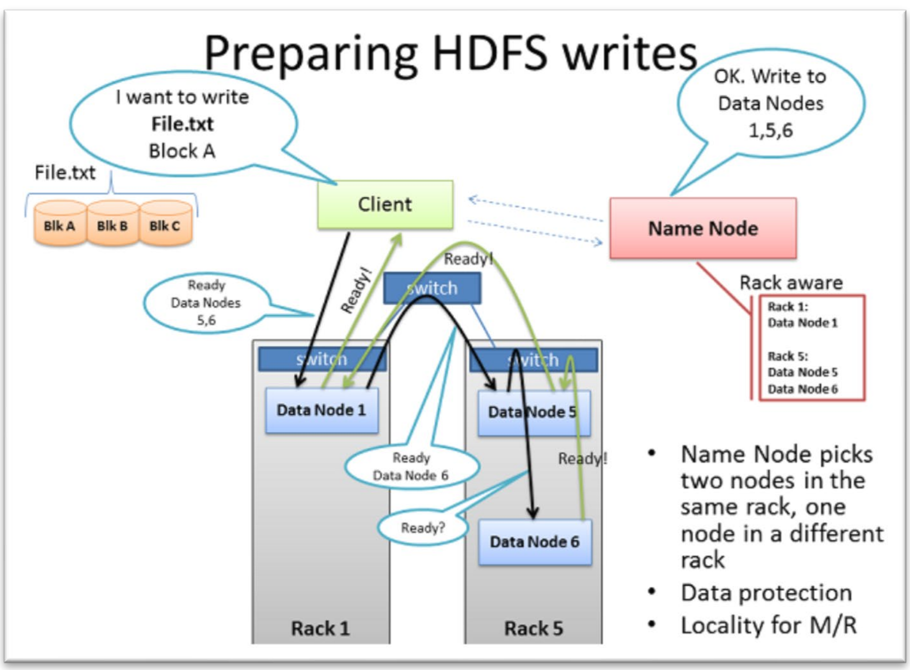

77. The Accused Instrumentalities include network switches.

78. 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.

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

80. 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.

81. 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.

82. 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.

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

84. 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.

85. 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.

86. 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.

87. Uber is on notice of the infringing products, services, features, and how Uber operates the Accused Instrumentalities to perform the claimed methods and use the claimed apparatuses.

88. Uber's infringing conduct has damaged MNS.

89. Uber is 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

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

91. Uber has infringed and continues to infringe at least claims 1 and 2 of the '388 Patent by making, using, selling, and/or offering to sell the Accused Instrumentalities (the Uber Big Data Platform).

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

93. 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.

94. 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.

95. 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.

96. 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).

97. 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).

98. 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.

99. 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).

100. Uber is on notice of the infringing products, services, features, and how Uber operates the Accused Instrumentalities to perform the claimed methods and use the claimed apparatuses.

101. Uber's infringing conduct has damaged MNS.

102. Uber is liable to MNS in an amount that adequately compensates it for Uber's 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 Uber has 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 Uber's infringement up until the date such judgment is entered, including prejudgment and postjudgment interest, costs, and disbursements as justified under 35 U.S.C. § 284 and an accounting, if necessary to adequately compensate MNS for Uber's 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 10, 2019

By /s/ *Stephen M. Lobbin*

Cabrach J. Connor
Texas Bar No. 24036390 (*pro hac vice forthcoming*)
cab@connorkudlaclee.com
Kevin S. Kudlac
Texas Bar No. 00790089 (*pro hac vice forthcoming*)
kevin@connorkudlaclee.com
Jennifer Tatum Lee
Texas Bar No. 24046950 (*pro hac vice forthcoming*)
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 Network Solutions, LLC