

**UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF TENNESSEE**

**UNIVERSITY OF TENNESSEE RESEARCH
FOUNDATION AND SAINT MATTHEW
RESEARCH, LLC,**

Plaintiffs,

v.

**AMAZON.COM, INC. AND AMAZON WEB
SERVICES, INC.**

Defendants.

Civil Action No. _____

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

The University of Tennessee Research Foundation and Saint Matthew Research, LLC (collectively, “Plaintiffs”), by its undersigned counsel, bring this action and make the following allegations of patent infringement relating to U.S. Patent Nos.: 7,454,411 (the “411 patent”); 8,099,733 (the “733 patent”); 6,741,983 (the “983 patent”); 7,272,612 (the “612 patent”); and 7,882,106 (the “106 patent”) (collectively, the “patents-in-suit”). Defendants Amazon.com, Inc. (“Amazon.com”) and Amazon Web Services, Inc. (“AWS”) (collectively, “Amazon” or “Defendant”) infringes each of the patents-in-suit in violation of the patent laws of the United States of America, 35 U.S.C. § 1 *et seq.*

INTRODUCTION

1. This is an action for infringement of patents awarded to computer scientists for their work at The University of Tennessee. The patents are owned by the University of Tennessee Research Foundation, a non-profit 501(c)(3) organization that promotes, licenses, and commercializes The University of Tennessee’s intellectual property.

2. Originally created in 1935 as the University of Tennessee Research Corporation, the University of Tennessee Research Foundation helps The University of Tennessee fulfill its

mission in becoming a national leader in research, discovery, and innovation. The University of Tennessee Research Foundation “was established to protect, manage, and commercialize university inventions and intellectual property; grow the university research enterprise; develop and support an entrepreneurial culture; and contribute to state and regional economic development.”¹

3. This case arises from Amazon’s infringement of a portfolio of patents protecting the groundbreaking work of several of The University of Tennessee’s faculty in the fields of parallel processing and high performance database design. The patents disclose systems and methods for parallel processing computer architectures that evenly distribute computational loads over multiple nodes. The patents also teach using the structural properties of data in large datasets to improve database performance. The inventions taught in the patents employ parallel processing schemes, which drastically improve the storage and retrieval capabilities of databases.

4. Established in 1794, The University of Tennessee is Tennessee’s flagship public research institution based in Knoxville, Tennessee. The University of Tennessee has more than 28,000 students, 1700 faculty members, and over 300 degree programs. In 2016, The University of Tennessee system had \$435 million in research and sponsored program expenditures.² To maximize the public benefit that its research generates, in certain instances, technological innovations developed by The University of Tennessee faculty, are assigned to the University of Tennessee Research Foundation, which patents and commercializes these innovations on behalf of The University of Tennessee.

¹ *State of Tennessee Comptroller of the Treasury, THE UNIVERSITY OF TENNESSEE FINANCIAL AND COMPLIANCE AUDIT REPORT* at 19 and 57 (January 20, 2015).

² Joe DiPietro, *THE STATE OF THE UNIVERSITY OF TENNESSEE ADDRESS* (February 21, 2017), *available at*: <http://tennessee.edu/state-of-ut/>.

5. In the late 1990's Professor J. Douglas Birdwell and his colleagues at The University of Tennessee developed technologies for searching large volumes of DNA profile data. Initially, Professor Birdwell and his fellow researchers sought to design database systems that could handle the vast and multi-dimensional nature of datasets associated with storing DNA profiles. DNA profile information presented unique challenges because of the need to search through tens of millions of records using complex match specifications. At the time, standard database structures failed to provide "rapid access to records" and were unable to "take advantage of naturally occurring structure in the data."³

6. To handle large datasets of DNA profile information, The University of Tennessee researchers developed technologies for efficiently distributing computationally intensive query functionality between computer hosts, grouping data records into distinct clusters in a database, and coordinating communication between groups of computer processors. Professor Birdwell and his colleagues developed technologies that used the structural properties of DNA data to improve the retrieval of data records. In addition, The University of Tennessee researchers developed parallel processing computer architectures for balancing computational requests across nodes. "As Search Queues become unbalanced, unprocessed Search Requests are exchanged to bring them back into balance. This exchange occurs randomly with a stochastic selection method utilized to determine the recipient of each exchange."⁴

7. While Dr. Birdwell and his collaborators initially sought to address the need for database systems capable of handling the size and complex nature of DNA profile information, their work led to groundbreaking innovations applicable to database systems that handle a wide

³ '983 patent, Col. 1:48-51.

⁴ *Id.* at Col. 28:2-6.

variety of complex and large datasets. The technologies developed for handling large volumes of DNA profile data were identified by The University of Tennessee researchers as providing groundbreaking insights applicable far beyond forensic applications.

8. Highlighting the importance of The University of Tennessee faculty's contributions to the field is the fact that the patents-in-suit have been cited by over 300 U.S. Patents and Patent Applications by a wide variety of the largest companies operating in the field. For example, the patents-in-suit have been cited by companies such as:

- International Business Machines Corporation (IBM)⁵
- Hewlett-Packard Enterprise Company (HPE)⁶
- Google, Inc.⁷
- Oracle Corporation⁸
- SAP SE⁹
- Microsoft Corporation¹⁰
- EMC Corporation¹¹
- Qualcomm, Inc.¹²
- Koninklijke Philips N.V. (Philips)¹³
- Western Digital Corporation¹⁴
- Teradata Corp.¹⁵
- Xerox Corporation¹⁶

⁵ See, e.g., U.S. Patent Nos. 6,915,289; 7,269,786; 7,684,963; 7,685,584.

⁶ See, e.g., U.S. Patent Nos. 8,621,480; 8,185,893; 8,341,626.

⁷ See, e.g., U.S. Patent Nos. 8,326,861; 8,316,019; 8,392,396; 8,73,238.

⁸ See, e.g., U.S. Patent Nos. 7,174,344; 7,779,008; 8,229,909; 9,436,514.

⁹ See, e.g., U.S. Patent Nos. 8,806,016; 8,103,772; 8,499,078; 9,444,732.

¹⁰ See, e.g., U.S. Patent Nos. 7,921,424; 7,502,807; 9,430,508.

¹¹ See, e.g., U.S. Patent Nos. 9,171,042; 9,275,117; 9,454,573.

¹² See, e.g., U.S. Patent No. 9,110,726.

¹³ See, e.g., U.S. Patent App. No. 14/416,647.

¹⁴ See, e.g., U.S. Patent Nos. 8,868,487; 8,874,515; 8,954,385.

¹⁵ See, e.g., U.S. Patent Nos. 7,359,913; 8,938,444.

¹⁶ See, e.g., U.S. Patent No. 7,720,848.

THE UNIVERSITY OF TENNESSEE'S LANDMARK PARALLEL PROCESSING TECHNOLOGIES

9. Through their work at The University of Tennessee, Professor Birdwell and his colleagues developed groundbreaking inventions in the fields of parallel processing and high performance database design

10. Throughout the 2000's, these developments, which revolutionized the scalability of modern-day database systems, were widely adopted in the market. Dr. Birdwell recognized in a subsequent academic paper that the groundbreaking patents came out of a need to develop technologies for handling large DNA data sets.

The projected growth of the NDIS [National DNA Index System] database and the demand for searches of the database necessitates migration to a parallel computing platform. Effective utilization of a parallel computer architecture requires the computational load to be distributed more or less evenly over the available CEs. The qualifier “more or less” is used because the communications required to distribute the load consume both computational resources and network bandwidth. A point of diminishing returns exists. The distribution of computational load across available resources is referred to as the load balancing problem in the literature.

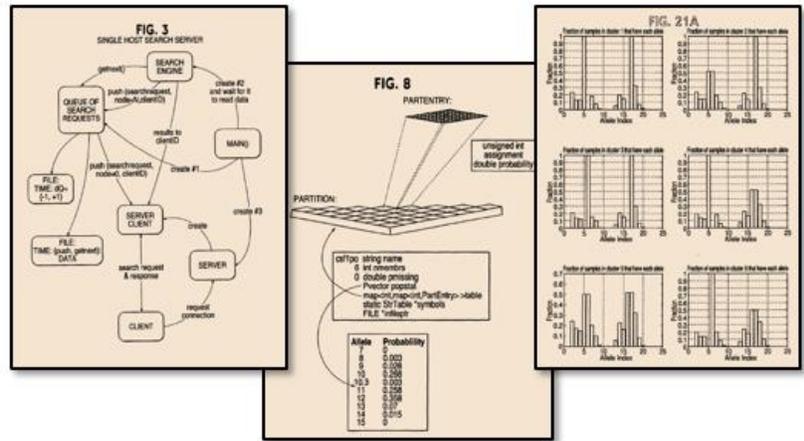
Douglas J. Birdwell et al, *Dynamic Time Delay Models for Load Balancing Part I: Deterministic Models*, in ADVANCES IN TIME-DELAY SYSTEMS, pp. 355-370. Springer Berlin Heidelberg (2004) (emphasis added).

11. In the years following the development of the patents-in-suit, the volume of data has grown exponentially, requiring new technologies such as those disclosed in the patents-in-suit. Many of the companies that cite the patents-in-suit have also acknowledged that the growing volume of data led to unique challenges that required new computer technologies.

With the surging volume of data being sourced from an ever-growing variety of data sources and applications, many streaming with great velocity, organizations are unable to use traditional data integration mechanisms such as ETL (extraction, transformation, and load). Big Data requires new strategies and technologies designed to analyze big data sets at terabyte or even petabyte scale.

An Enterprise Architect's Guide to Big Data – Reference Architecture Overview, ORACLE ENTERPRISE ARCHITECTURE WHITE PAPER at 26 (March 2016).

12. Each of the patents-in-suit claim priority to 1999, when typical datasets stored in databases were orders of magnitude smaller than today. Dr. Birdwell and his colleagues invented technologies that today are incorporated into the products of many leading technology companies.



‘983 Patent, figs. 3, 8 and 21A (Figures from the patent showing a “schematic representation of a search server residing on a single host computer,” a “schematic representation of the C++ Partition object,” and a “scores cluster.”).

13. The inventors of the patents-in-suit include distinguished members of The University of Tennessee faculty. The inventors are:

14. Professor J. Douglas Birdwell is a Professor Emeritus of Electrical and Computer Engineering at The University of Tennessee, Knoxville. Dr. Birdwell’s research into computing and information systems, including high-performance databases, spans from the 1970’s through today. Dr. Birdwell is the author of over one hundred publications and his research areas focus on control systems, computer hardware and software applications, signal processing and artificial intelligence, and intelligent process supervision. Dr. Birdwell has received research grants from

the National Science Foundation,¹⁷ U.S. Department of Justice,¹⁸ and NASA.¹⁹

15. Dr. Birdwell is a Fellow of the Institute of Electrical and Electronics Engineers (“IEEE”), and in December 2015, Dr. Birdwell was named a Fellow of the National Academy of Inventors.²⁰ Dr. Birdwell received his B.S. and M.S. degrees in Electrical Engineering from The University of Tennessee, and in 1978 he received his Ph.D. in Electrical Engineering from the Massachusetts Institute of Technology (“MIT”). Upon receiving his doctorate, Dr. Birdwell joined The University of Tennessee faculty in 1978.

16. Professor Tsewei Wang is a Professor Emeritus of Chemical and Biomolecular Engineering at The University of Tennessee, Knoxville. Dr. Wang is the author of numerous articles on load balancing, data mining, and clustering data records. Dr. Wang received her Master’s degree in Chemical Engineering from The University of Tennessee and her Ph.D. in Biophysics from MIT. Dr. Wang’s research focuses on the fields of data mining, process monitoring, and bioinformatics.

17. Professor Roger Horn is a Research Professor of Electrical Engineering and Computer Science at The University of Tennessee, Knoxville. Dr. Horn’s research and industrial experience range from digital circuit design, data acquisition, control systems, and the development of application and driver software to turbine control systems, radiation monitoring systems, smart pressure transducers, and medical devices. Dr. Horn is a Registered Professional

¹⁷ NATIONAL SCIENCE FOUNDATION GRANT NO. ANI-0312611 (2003) (“[F]or modeling, optimization and testing of an innovative load balancing strategies in large-scale, distributed-computing systems consisting of geographically-distant computational elements.”).

¹⁸ U.S. DEPARTMENT OF JUSTICE GRANT NO. J-FBI-98-083.

¹⁹ NASA-AMES RESEARCH CENTER GRANT NO. NGL-22-009-124.

²⁰ *Third UT Faculty Member Named to National Academy of Inventors*, TENNESSEE TODAY (Dec. 17, 2015), available at, <http://tntoday.utk.edu/2015/12/17/birdwell-ut-named-national-academy-inventors/>.

Engineer in the State of Tennessee, and a member of the National Society of Professional Engineers, the Institute of Electrical and Electronics Engineers, the IEEE Control Systems Society, and the IEEE Robotics and Automation Society. Dr. Horn received his Bachelor's and Master's degrees in Electrical Engineering from the University of California, Berkeley in 1978 and 1984, respectively. Dr. Horn received his Ph.D. in Electrical Engineering from The University of Tennessee in 1992.

18. Professor David J. Icové is a UL Professor of Electrical Engineering and Computer Science at The University of Tennessee, Knoxville. Dr. Icové is the author of numerous articles and textbooks on forensic analysis, computational modeling, and pattern recognition. Dr. Icové received his Master's in Electrical Engineering from The University of Tennessee in 1973, and his Ph.D. in Engineering Science and Mechanics from The University of Tennessee in 1979. Dr. Icové is the inaugural Underwriters Laboratories Professor of Practice at The University of Tennessee.²¹

19. Puneet Yadav is a Director of Product Marketing at Lam Research in Fremont, California. Mr. Yadav received his Bachelor's degree in Chemical Engineering from the Indian Institute of Technology in 1998 and a Master's degree in Chemical Engineering from The University of Tennessee, Knoxville in 2001. At the time the inventions disclosed in the patents-in-suit were developed, Mr. Yadav was a graduate student at The University of Tennessee and collaborated with the other inventors of the patents-in-suit.²²

²¹ *Underwriters Laboratories Supporting College of Engineering Course*, TENNESSEE TODAY (April 7, 2014), available at, <http://tntoday.utk.edu/2014/04/07/underwriters-laboratories-supports-ut-engineering-course/>.

²² Puneet Yadav, *Increasing the Speed and Efficiency of Search in FBI/CODIS DNA Database Through Multivariate Statistical Clustering Approach and Development of a Similarity Ranking Scheme*, MASTER'S THESIS UNIVERSITY OF TENNESSEE (2001).

THE PARTIES

UNIVERSITY OF TENNESSEE RESEARCH FOUNDATION

20. The University of Tennessee Research Foundation is the second-oldest university research foundation in the United States. The University of Tennessee Research Foundation's mission is to promote, support, and carry out the research mission of The University of Tennessee, to enhance the competitive position of The University of Tennessee for research and development funding, facilitate expanded research and development activities at The University of Tennessee, and to facilitate the commercialization of University of Tennessee research outcomes and the transfer of research-generated technology from The University of Tennessee to commercial and industrial enterprises in furtherance of the economic development of the State of Tennessee.

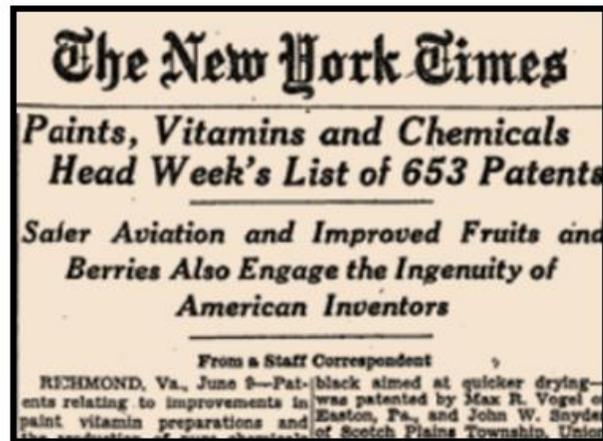
21. UTRF is responsible for licensing patents covering University of Tennessee inventions and collecting royalties on behalf of The University of Tennessee from those license agreements. The State of Tennessee Comptroller has described the role of UTRF as:

The foundation's stated purpose is, in conjunction with the university, to grow the University of Tennessee research enterprise; harvest, manage, and market University of Tennessee intellectual property; encourage and support entrepreneurial education and ventures by faculty, staff, students, and commercial partners/affiliates of the University of Tennessee; and to contribute to the well-being of the State of Tennessee through economic development.

The University of Tennessee Audit Report for The Year Ending June 30, 2014, STATE OF TENNESSEE COMPTROLLER OF THE TREASURY at 57 (January 20, 2015).

22. Since its inception as the University of Tennessee Research Corporation in 1935, UTRF has been actively involved in licensing patents whose technologies were first developed at The University of Tennessee. The University of Tennessee Research Corporation's 1935 charter stated that its mission was "to promote, encourage and aid scientific social and/or educational

investigation and research.”²³ In furtherance of these objectives, the corporation was empowered “[t]o aid in the prosecution of applications for patents, registrations and/or copyrights, foreign and domestic To prosecute infringements or invasions of any patent, trade-mark, trade name, brand, label, copyright or patent right in which the corporation may be interested.”²⁴



Safer Aviation and Improved Fruits and Berries Also Engage the Ingenuity of American Inventors, N.Y. TIMES at 20 (June 10, 1944) (Describing a patent “issued to Brooks D. Drain of Knoxville, Tenn., on a ‘healthy, prolific strawberry plant variety’” and “assigned to the University of Tennessee Research Corporation.”).

23. In its first decades of existence the University of Tennessee Research Foundation licensed its patents to cottonseed, fruit varieties, and water treatment.

Of late, research on a small commercial scale is being carried on by the University of Tennessee Experiment Station and a *University of Tennessee Research Corporation* has been set up to license and control the use of the process. Here then we have a state institution of learning working with a federal government agency, a professional society, a trade association, and private capital all working together for a common purpose now by the upbuilding of regional agricultural economy.”²⁵

The largest manufacturer of cottonseed manufacturing equipment in the United States has applied for a license under royalty to be paid to the University of

²³ CHARTER OF INCORPORATION OF THE UNIVERSITY OF TENNESSEE RESEARCH CORPORATION (December 31, 1934).

²⁴ *Id.*

²⁵ THE UNIVERSITY OF TENNESSEE RECORD, Vol. 44, Issue 4 (1941).

Tennessee Research Corporation.²⁶

The University of Tennessee Research Corporation licenses manufacturers to make equipment resulting from the research work for which a patent has been obtained or applied for.”²⁷

24. More recently, UTRF has licensed patents pertaining to various fields of technology, including important, lifesaving technologies. UTRF has recently licensed patents for a vaccine for group A streptococcus (Strep-A) and for treatments for prostate cancer.²⁸

SAINT MATTHEW RESEARCH, LLC

25. Saint Matthew Research, LLC (“SMR”) is a California limited liability company with its principal place of business at 11400 W. Olympic Blvd., Suite 200, Los Angeles, California 90064. Pursuant to a license agreement with UTRF, SMR is the exclusive licensee of the patents-in-suit. SMR facilitates the University of Tennessee Research Foundation’s mission of commercializing UT’s technologies by assisting UTRF’s commercialization efforts.

AMAZON

26. On information and belief, Amazon.com, Inc. is a Delaware corporation with a principal place of business at 410 Terry Avenue North, Seattle, WA 98109. Amazon may be served with process by delivering a summons and a true and correct copy of this complaint to its registered agent for receipt of service of process, Corporation Service Company, 2711

²⁶ UNITED STATES CONGRESS JOINT COMMITTEE HEARING ON THE TENNESSEE VALLEY AUTHORITY at 1451 (1939).

²⁷ *Statement of David E. Lilienthal, Chairman, Tennessee Valley Authority, TECHNOLOGICAL MOBILIZATION, HEARINGS BEFORE A SUBCOMMITTEE OF THE COMMITTEE ON MILITARY AFFAIRS OF THE UNITED STATES SENATE S. 2721 Vol. 3 at 910 (December 1942).*

²⁸ James B. Dale et al., *Potential Coverage of a Multivalent M Protein-Based Group a Streptococcal Vaccine*, VACCINE 31.12 (2013): 1576–1581 (“The University of Tennessee Research Corporation has licensed the technology to Vaxent, LLC.”); Tom Wilemon, *Powerful Possibilities: GTx Scientists Battle Time, Regulatory Setbacks to Make it Big*, MEMPHIS DAILY NEWS (May 17, 2010) (“The SARM patents, as well as Steiner’s research into tormifene for the prevention of prostate cancer, were licensed to GTx by the University of Tennessee.”).

Centerville Road, Suite 400, Wilmington, Delaware 19808.

27. On information and belief, Amazon Web Services, Inc. is a Delaware corporation with its principal office at 410 Terry Avenue North, Seattle, WA 98109. Amazon Web Services can be served through its registered agent, Corporation Service Company, 2711 Centerville Road, Suite 400, Wilmington, Delaware 19808.

28. On information and belief, Amazon has offices in the State of Tennessee where it sells, develops, and/or markets its products including offices in Chattanooga, Charleston, Lebanon, and Nashville.²⁹

JURISDICTION AND VENUE

29. This action arises under the patent laws of the United States, Title 35 of the United States Code. Accordingly, this Court has exclusive subject matter jurisdiction over this action under 28 U.S.C. §§ 1331 and 1338(a).

30. Upon information and belief, this Court has personal jurisdiction over Amazon in this action because Amazon has committed acts within the Eastern District of Tennessee giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Amazon would not offend traditional notions of fair play and substantial justice. Defendant Amazon, directly and/or through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the patents-in-suit. Moreover, Amazon actively directs its activities to customers located in the State of Tennessee.

²⁹ *Amazon Jobs*, AMAZON WEBSITE, available at: <https://www.amazon.jobs/en> (last visited April 2017).

31. Venue is proper in this district under 28 U.S.C. §§ 1391(b)-(d) and 1400(b).

Defendant Amazon, upon information and belief, has transacted business in the Eastern District of Tennessee and has committed acts of direct and indirect infringement in the Eastern District of Tennessee.

TECHNOLOGY BACKGROUND

U.S. PATENT NO. 7,454,411

32. U.S. Patent No. 7,454,411 (“the ‘411 patent”) entitled, *Parallel Data Processing Architecture*, was filed on January 30, 2004, and claims priority to September 28, 1999. The ‘411 patent is subject to a 35 U.S.C. § 154(b) term extension of 385 days. UTRF is the owner by assignment of the ‘411 patent. SMR is the exclusive licensee of the ‘411 patent. A true and correct copy of the ‘411 patent is attached hereto as Exhibit A.

33. The ‘411 patent teaches a novel parallel data processing system for search, storage, and retrieval of data from a database.

34. The ‘411 patent and its underlying patent application have been cited by 83 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ‘411 patent and its underlying patent application as relevant prior art:

- International Business Machines Corporation
- EMC Corporation
- AT&T, Inc.
- Koutin Technology Inc.
- Hewlett-Packard Enterprise Company
- Google, Inc.
- Infosys Ltd.
- Oracle Corporation
- Qualcomm, Inc.
- Raytheon Company
- Sandbridge Technologies, Inc.

U.S. PATENT NO. 8,099,733

35. U.S. Patent No. 8,099,733 (“the ‘733 patent”) entitled, *Parallel Data Processing Architecture*, was filed on January 2, 2008, and claims priority to September 28, 1999. The ‘733 patent is subject to a 35 U.S.C. § 154(b) term extension of 1,051 days. UTRF is the owner by assignment of the ‘733 patent. SMR is the exclusive licensee of the ‘733 patent. A true and correct copy of the ‘733 patent is attached hereto as Exhibit B.

36. The ‘733 patent teaches a parallel data processing architecture containing two or more host processors. One of the host processors responds to client queries for database records. The host processors are further able to communicate with one another via a communication systems. The parallel data processing architecture balances the workload between the host processors.

37. The ‘733 patent teaches the balancing of a workload between host processors through the host processors broadcasting load information (e.g., processor capacity and the length of the search queue). The load information is broadcasted by host processors to at least one or more host processors in the system.

38. The parallel data processing architecture disclosed in the ‘733 patent is capable of balancing the workload between two or more host processors by having the host processors bring their associated search queue of client requests into balance with at least one other host processor. The workload balancing by the host processors use a time constant responsive to receipt of broadcasted capacity and load information.

39. The ‘733 patent describes balancing workload between host processors as being enabled through the exchange of block search requests between two or more of the host processors. This exchange of search requests between host processors minimizes the time required to respond to client queries by adjusting the block of search requests that are exchanged

based on the relative processing speeds of host processors and the communications protocol between host processors.

40. The '733 patent balances workload between computer hosts to equalize the average wait time for computation of requests. For example, blocks of search requests can be exchanged among hosts from hosts with relatively long average waiting times to hosts with shorter waiting times. The sizes of the blocks of exchanged requests can be adjusted to accommodate the relative speeds of the processors and the inter-processor communications protocols and hardware.

41. Another insight for improving the performance of requests on a database system described by the '733 patent is to have each host broadcast measurements of their current load (queue lengths) and measurements of their capacity.

42. The '733 patent and its underlying patent application has been cited by 14 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '733 patent and its underlying patent application as relevant prior art:

- Oxford University Innovation Ltd.
- San Diego State University Foundation
- Path, Inc.
- Four J's Development Tools, Inc.
- Hewlett-Packard Enterprise Company

U.S. PATENT NO. 6,741,983

43. U.S. Patent No. 6,741,983 (“the ‘983 patent”) entitled, *Method of Indexed Storage and Retrieval of Multidimensional Information*, was filed on September 28, 2000, and claims priority to September 28, 1999. The ‘983 patent is subject to a 35 U.S.C. § 154(b) term extension of 652 days. UTRF is the owner by assignment of the ‘983 patent. SMR is the

exclusive licensee of the '983 patent. A true and correct copy of the '983 patent is attached hereto as Exhibit C.

44. The '983 patent teaches a method of representing one or more variables in each data record in a binary form, whereby the value of each bit is assigned based on the value of a variable.

45. The '983 patent further teaches the use of principal component analysis to the chosen set of variables for selected data records where at least two principal component vectors are identified.

46. The '983 patent further teaches organizing data records in a database into clusters to choose a set of variables from the data records where the principal component analysis of the set of variables yields distinct clusters of data records.

47. The '983 patent has been cited by 85 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '983 patent as relevant prior art:

- International Business Machines Corporation
- Oracle Corporation
- SAP SE
- Hewlett-Packard Enterprise Company
- Google, Inc.
- Honeywell International, Inc.
- Business Objects Software Limited
- Infosys, Ltd.
- NC, Corporation
- Koutin Technology, Inc.
- Koninklijke Philips N.V.
- Sony Corporation
- Microsoft Corporation
- Teradata Corporation

U.S. PATENT NO. 7,272,612

48. U.S. Patent No. 7,272,612 (“the ‘612 patent”) entitled, *Method of Partitioning Data Records*, was filed on January 30, 2004, and claims priority to September 28, 1999. The ‘612 patent is subject to a 35 U.S.C. § 154(b) term extension of 498 days. UTRF is the owner by assignment of the ‘612 patent. SMR is the exclusive licensee of the ‘612 patent. A true and correct copy of the ‘612 patent is attached hereto as Exhibit D.

49. The ‘612 patent teaches defining a function of a distribution of values of a designated variable associated with the data records wherein the function comprises a combination of measures of entropy and adjacency. Further, the ‘612 patent teaches weighting the adjacency by a weighting factor.

50. The ‘612 patent teaches partitioning computer data records by partitioning the values of the designated variable into two or more groups. The value of the function is determined by applying an optimization procedure and assigning a data record to a group according to the values of the designated variable.

51. The ‘612 patent and its underlying patent application have been cited by 64 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ‘612 patent and its underlying patent application as relevant prior art:

- Yahoo! Inc.
- Alcatel-Lucent SA
- AT&T, Inc.
- CA, Inc.
- International Business Machines Corporation
- Google, Inc.
- Nuance Communications, Inc.
- Oracle Corporation
- Koutin Technology Inc. Ltd.
- Infosys Ltd.

- SAP SE
- Microsoft Corporation
- Agilent Technologies, Inc.
- Qualcomm, Inc.
- Sandbridge Technologies, Inc.
- Sony Corporation
- Telefonaktiebolaget L M Ericsson

U.S. PATENT NO. 7,882,106

52. U.S. Patent No. 7,882,106 (“the ‘106 patent”) entitled, *Method of Indexed Storage and Retrieval of Multidimensional Information*, was filed on August 10, 2007, and claims priority to September 28, 1999. The ‘106 patent is subject to a 35 U.S.C. § 154(b) term extension of 844 days. UTRF is the owner by assignment of the ‘106 patent. SMR is the exclusive licensee of the ‘106 patent. A true and correct copy of the ‘106 patent is attached hereto as Exhibit E.

53. The inventions taught in the ‘106 patent achieve improvements in computer database systems by exploiting the natural structure of data records. Implementation of the system and methods disclosed in the ‘106 patent map to parallel computer architectures allowing scaling up for large volumes of database records.

54. The ‘106 patent teaches the partitioning of data records in a database by defining a function of a distribution of values of a designated variable associated with the database records. The function of a distribution of values of the designated variable comprises a combination of measures of entropy and adjacency. The measure of adjacency that is incorporated into the function is weighted by a weighting factor.

55. The ‘106 patent discloses partitioning database records by partitioning the values of the designated variable. The partitioned database records are then stored in computer processor memory.

56. The '106 patent and its underlying patent application have been cited by 17 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '106 patent and its underlying patent application as relevant prior art:

- Adobe Systems, Inc.
- EMC Corporation
- West Digital Corporation
- SanDisk Corporation
- President and Fellows of Harvard College
- NuWave Technologies, Inc.
- Xerox Corporation

COUNT I
INFRINGEMENT OF U.S. PATENT NO. 7,454,411

57. Plaintiffs reference and incorporate by reference the preceding paragraphs of this Complaint as if fully set forth herein.

58. Amazon designs, makes, uses, sells, and/or offers for sale in the United States products and/or services for processing data within a database whereby tests based upon selected vectors are performed on each data record and the data records organized into clusters.

59. Amazon designs, makes, sells, offers to sell, imports, and/or uses Amazon Redshift (including functionality for: Amazon High Performance Computing ("HPC), Amazon CloudWatch, Amazon Auto Scaling, and Amazon Elastic Load Balancing functionality) (collectively, the "Amazon '411 Product(s)").

60. On information and belief, one or more Amazon subsidiaries and/or affiliates use the Amazon '411 Products in regular business operations.

61. On information and belief, one or more of the Amazon '411 Products are parallel data processing systems for search, storage, and retrieval of data of a database responsive to client queries for specific data from a database.

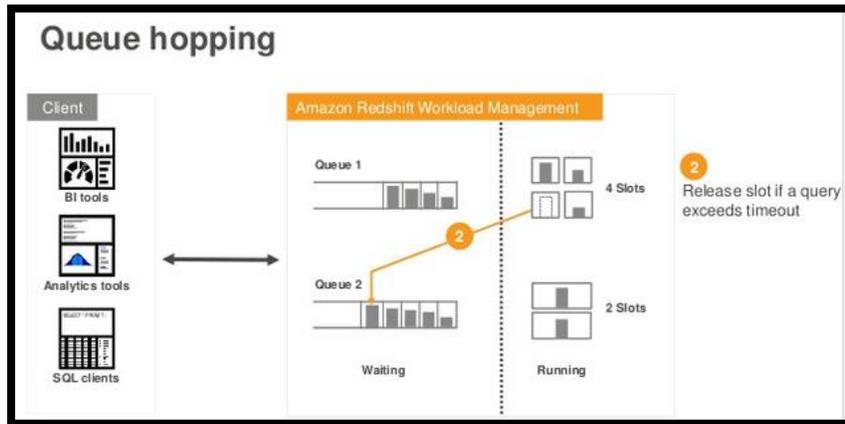
database wherein at least two host processors have a search engine and maintain information of a search queue of the client queries.

64. On information and belief, the Amazon '411 Products include technology comprising at least two host processors having a queue of search requests for specific data of a database, wherein each host processor executes a search engine, communicating capacity and load information between host processors and the at least two host processors exchanging at least one search request.

65. On information and belief, the Amazon '411 Products include search engines removing at least one search request from a search queue and generating an additional search request.

66. On information and belief, the Amazon '411 Products include host and root host processors maintaining a list of available host processors and information about the capacity and load for each available host processor in memory and broadcasts its capacity and load information to the other host processors and bring its search queue into balance with another host processor according to a time constant in response to receipt of the broadcast capacity and load information.

67. On information and belief, the Amazon '411 Products include a communications system coupling the host and root processors, wherein at least two host processors communicate capacity and load information to other host processors.



Maor Kleider, *Getting Started with Amazon Redshift*, AWS PRESENTATION at 19 (February 2017).

68. On information and belief, the Amazon ‘411 Products include technology allowing selected host processors to store a database index for the database comprising nodes of a database tree for the database and data accessible via the nodes of the database tree.

69. On information and belief, the Amazon ‘411 Products are available to businesses and individuals throughout the United States.

70. On information and belief, the Amazon ‘411 Products are provided to businesses and individuals located in the Eastern District of Tennessee.

71. On information and belief, Amazon has directly infringed and continues to directly infringe the ‘411 patent by, among other things, making, using, offering for sale, and/or selling technology for extracting data from sources of network-based information, including but not limited to the Amazon ‘411 Products, which include infringing technology for processing data within a database whereby tests based upon selected vectors are performed on each data record and the data records organized into clusters. Such products and/or services include, by way of example and without limitation, the Amazon ‘411 Products.

72. By making, using, testing, offering for sale, and/or selling products and services, including but not limited to the Amazon ‘411 Products, Amazon has injured Plaintiffs and is

liable to Plaintiffs for directly infringing one or more claims of the '411 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

73. On information and belief, Amazon also indirectly infringes the '411 patent by actively inducing infringement under 35 USC § 271(b).

74. On information and belief, Amazon has had knowledge of the '411 patent since at least September 17, 2014. On September 17, 2014, the University of Tennessee Research Foundation contacted Amazon via email in an attempt to secure a license to the patents-in-suit outside the context of litigation.

75. Alternatively, Amazon has had knowledge of the '411 patent since at least service of this Complaint or shortly thereafter, and on information and belief, Amazon knew of the '411 patent and knew of its infringement, including by way of this lawsuit.

76. On information and belief, Amazon intended to induce patent infringement by third-party customers and users of the Amazon '411 Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Amazon specifically intended and was aware that the normal and customary use of the accused products would infringe the '411 patent. Amazon performed the acts that constitute induced infringement, and would induce actual infringement, with knowledge of the '411 patent and with the knowledge that the induced acts would constitute infringement. For example, Amazon provides the Amazon '411 Products that have the capability of operating in a manner that infringe one or more of the claims of the '411 patent, including at least claim 1, and Amazon further provides documentation and training materials that cause customers and end users of the Amazon '411 Products to utilize the products in a

manner that directly infringe one or more claims of the '411 patent.³⁰ By providing instruction and training to customers and end-users on how to use the Amazon '411 Products in a manner that directly infringes one or more claims of the '411 patent, including at least claim 1, Amazon specifically intended to induce infringement of the '411 patent. On information and belief, Amazon engaged in such inducement to promote the sales of the Amazon '411 Products, e.g., through Amazon user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '411 patent. Accordingly, Amazon has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '411 patent, knowing that such use constitutes infringement of the '411 patent.

77. The '411 patent is well-known within the industry as demonstrated by multiple citations to the '411 patent in published patents and patent applications assigned to technology companies and academic institutions. Amazon is utilizing the technology claimed in the '411 patent without paying a reasonable royalty. Amazon is infringing the '411 patent in a manner

³⁰ See, e.g., Ben Butler, HIGH PERFORMANCE COMPUTING ON AWS, AMAZON WEB SERVICE PRESENTATION (2014); *Elastic Load Balancing User Guide*, AMAZON WEB SERVICES DOCUMENTATION (2017); Dan Brown, *Elastic Load Balancing: Deep Dive & Best Practices*, AWS RE: INVENT NET403 PRESENTATION (December 1, 2016); *Best Practices in Evaluating Elastic Load Balancing*, AMAZON WEB SERVICES WHITE PAPER (September 23, 2014); Alex Coqueiro, *High Performance Computing in AWS – Technical Workshop*, AMAZON WEST SERVICES PRESENTATION (2015); *Getting Maximum Performance from Amazon Redshift: High Concurrency, Fast Data Loads, and Complex Queries*, AWS RE:INVENT PRESENTATION (November 13, 2013); *Amazon Redshift Database Developer Guide API Version*, AWS DEVELOPER DOCUMENTATION (December 1, 2012); Tony Gibbs, *Amazon Redshift Deep Dive – Storage Subsystem and Query Life Cycle*, AWS PRESENTATION (February 2017); *Amazon Redshift Masterclass – Amazon Web Services*, AWS WEBINAR CHANNEL (October 28, 2015), available at: <https://www.youtube.com/watch?v=GgLKodmL5xE>; Maor Kleider, *Getting Started with Amazon Redshift*, AWS PRESENTATION (February 2017); and Ran Tessler, *Amazon Redshift Deep Dive*, AWS Presentation (July 23, 2015).

best described as willful, wanton, malicious, in bad faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate.

78. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '411 patent.

79. As a result of Amazon's infringement of the '411 patent, Plaintiffs have suffered monetary damages, and seek recovery in an amount adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty for the use made of the invention by Amazon together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 8,099,733

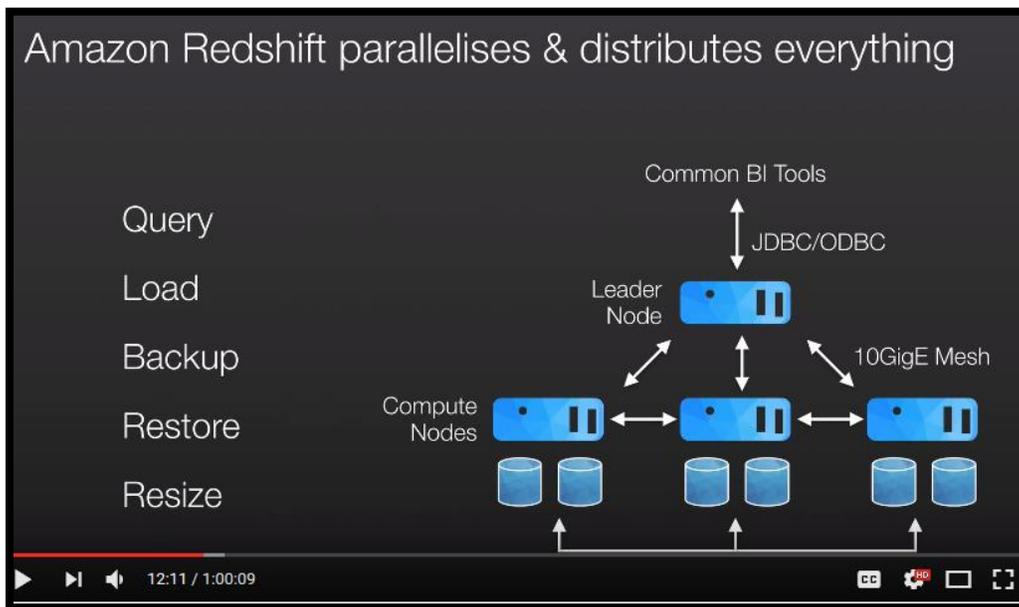
80. Plaintiffs reference and incorporate by reference the preceding paragraphs of this Complaint as if fully set forth herein.

81. Amazon designs, makes, uses, sells, and/or offers for sale in the United States products and/or services comprising a parallel data processing architecture for search, storage and retrieval of data of a database responsive to queries for specific data of a database where a plurality of host processors capable of communicating with one another and broadcasting load information (processor capacity and search queue length) and bringing client queries into balance with one another based on the receipt of the broadcasted capacity and load information.

82. Amazon designs, makes, sells, offers to sell, imports, and/or uses Amazon Redshift (including functionality for: Amazon High Performance Computing ("HPC"), Amazon CloudWatch, Amazon Auto Scaling, and Amazon Elastic Load Balancing functionality) (collectively, the "Amazon '733 Product(s)").

83. On information and belief, one or more Amazon subsidiaries and/or affiliates use the Amazon '733 Products in regular business operations.

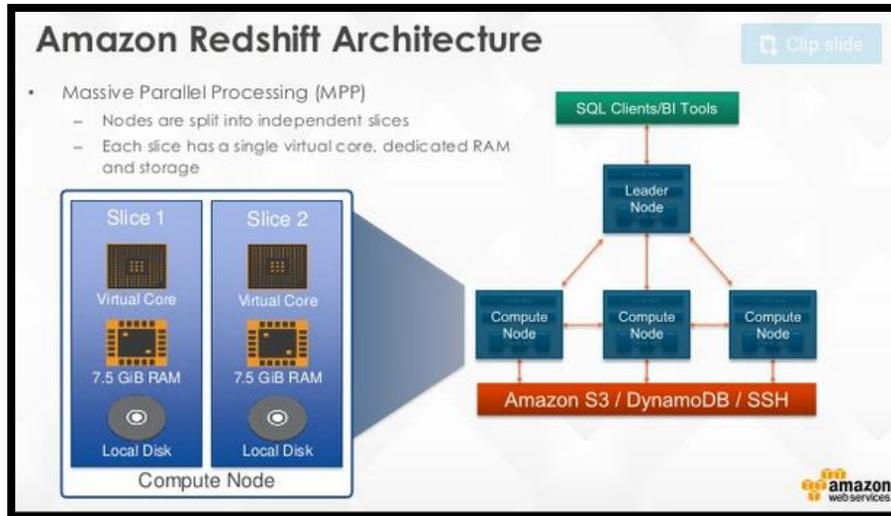
84. On information and belief, one or more of the Amazon ‘733 Products comprise a parallel data processing architecture for search, storage and retrieval of data of a database responsive to queries for specific data of said database. Specifically, the Amazon ‘733 Products comprise systems where there are a plurality of host processors and where at least one root host processor is responsive to a client query for data from a database.



Amazon Redshift Masterclass – Amazon Web Services, AWS WEBINAR CHANNEL at 12:11 (October 28, 2015), available at: <https://www.youtube.com/watch?v=GgLKodmL5xE>.

85. On information and belief, one or more of the Amazon ‘733 Products comprise a plurality of host processors comprising at least one root host processor responsive to a client query for specific data of a database and at least one other host processor.

86. On information and belief, one or more of the Amazon ‘733 Products comprise a communication system and method for coupling a plurality of host processors that communicate with one another their capacity and measurements of current loads.



Ran Tessler, *Amazon Redshift Deep Dive*, AWS Presentation at 3 (July 23, 2015).

87. On information and belief, one or more of the Amazon ‘733 Products comprise a method for balancing workload between said plurality of host processors.

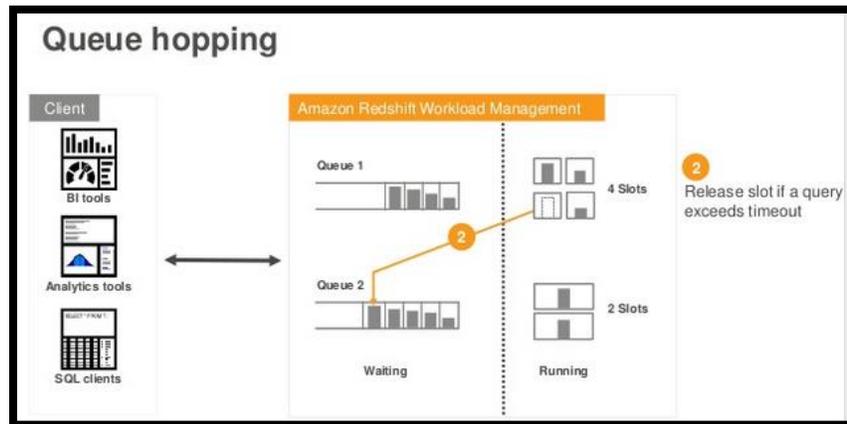
88. On information and belief, one or more of the Amazon ‘733 Products maintain load information of processor capacity and search queue length for each host processor.

89. On information and belief, one or more of the Amazon ‘733 Products maintain information of a search queue of client queries at the host processor for specific data within a database.

90. The Amazon ‘733 Products are systems which contain functionality for a plurality of host processors maintaining load information of processor capacity and search queue length of a host processor. The Amazon ‘733 Products enable each host to maintain information regarding both the capacity of the host and the current load on that host.

91. On information and belief, one or more of the Amazon ‘733 Products enable each host processor to broadcast load information of its processor capacity and search queue length to at least one other of a plurality of host processors.

92. On information and belief, one or more of the Amazon ‘733 Products enable each of the host processors to reconfigure information on available host processors responsive to the receipt of broadcast information of capacity and load information.



Maor Kleider, *Getting Started with Amazon Redshift*, AWS PRESENTATION at 19 (February 2017).

93. On information and belief, one or more of the Amazon ‘733 Products enable each host processor to bring its search queue of client queries into balance with another of a plurality of host processors according to a time constant responsive to receipt of the broadcast capacity and load information.

94. On information and belief, one or more of the Amazon ‘733 Products enable balancing search queues by exchanging unprocessed search requests with a recipient host processor responsive to a stochastic selection process to determine the recipient host processor of an exchanged search request between a root host processor and a recipient host processor.

95. The Amazon ‘733 Products balance exchanging unprocessed search requests with a recipient host processor responsive to a stochastic selection process to determine the recipient host processor of an exchanged search request between said root host processor and a recipient host processor.

96. On information and belief, one or more of the Amazon '733 Products enable exchanging a block of search requests between a plurality of host processors and adjusting a size of the block of exchanged requests according to relative processing speeds of host processors and inter-processor communications protocol between the host processors.

97. On information and belief, the Amazon '733 Products are available to businesses and individuals throughout the United States.

98. On information and belief, the Amazon '733 Products are provided to businesses and individuals located in the Eastern District of Tennessee.

99. On information and belief, Amazon has directly infringed and continues to directly infringe the '733 patent by, among other things, making, using, offering for sale, and/or selling technology for extracting data from sources of network-based information, including but not limited to the Amazon '733 Products, which include infringing technology comprising a parallel data processing architecture for search, storage, and retrieval of data of a database responsive to queries for specific data of said database where a plurality of host processors are capable of communicating with one another and broadcasting load information (processor capacity and search queue length) and bringing client queries into balance with one another based on the receipt of the broadcasted capacity and load information. Such products and/or services include, by way of example and without limitation, the Amazon '733 Products.

100. By making, using, testing, offering for sale, and/or selling products and services, including but not limited to the Amazon '733 Products, Amazon has injured Plaintiffs and is liable for directly infringing one or more claims of the '733 patent, including at least claim 18, pursuant to 35 U.S.C. § 271(a).

101. On information and belief, Amazon also indirectly infringes the ‘733 patent by actively inducing infringement under 35 USC § 271(b).

102. On information and belief, Amazon has had knowledge of the ‘733 patent since at least September 17, 2014. On September 17, 2014, the University of Tennessee Research Foundation contacted Amazon via email in an attempt to secure a license to the patents-in-suit outside the context of litigation.

103. Alternatively, Amazon has had knowledge of the ‘733 patent since at least service of this Complaint or shortly thereafter, and on information and belief, Amazon knew of the ‘733 patent and knew of its infringement, including by way of this lawsuit.

104. On information and belief, Amazon intended to induce patent infringement by third-party customers and users of the Amazon ‘733 Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Amazon specifically intended and was aware that the normal and customary use of the accused products would infringe the ‘733 patent. Amazon performed the acts that constitute induced infringement, and would induce actual infringement, with knowledge of the ‘733 patent and with the knowledge that the induced acts would constitute infringement. For example, Amazon provides the Amazon ‘733 Products that have the capability of operating in a manner that infringe one or more of the claims of the ‘733 patent, including at least claim 18, and Amazon further provides documentation and training materials that cause customers and end users of the Amazon ‘733 Products to utilize the products in a manner that directly infringe one or more claims of the ‘733 patent.³¹ By providing instruction

³¹ See, e.g., Ben Butler, HIGH PERFORMANCE COMPUTING ON AWS, AMAZON WEB SERVICE PRESENTATION (2014); *Elastic Load Balancing User Guide*, AMAZON WEB SERVICES DOCUMENTATION (2017); Dan Brown, *Elastic Load Balancing: Deep Dive & Best Practices*,

and training to customers and end-users on how to use the Amazon ‘733 Products in a manner that directly infringes one or more claims of the ‘733 patent, including at least claim 18, Amazon specifically intended to induce infringement of the ‘733 patent. On information and belief, Amazon engaged in such inducement to promote the sales of the Amazon ‘733 Products, e.g., through Amazon user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the ‘733 patent. Accordingly, Amazon has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ‘733 patent, knowing that such use constitutes infringement of the ‘733 patent.

105. The ‘733 patent is well-known within the industry as demonstrated by the multiple citations to the ‘733 patent in published patents and patent applications assigned to technology companies and academic institutions. Amazon is utilizing the technology claimed in the ‘733 patent without paying a reasonable royalty. Amazon is infringing the ‘733 patent in a manner best described as willful, wanton, malicious, in bad faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate.

AWS RE: INVENT NET403 PRESENTATION (December 1, 2016); *Best Practices in Evaluating Elastic Load Balancing*, AMAZON WEB SERVICES WHITE PAPER (September 23, 2014); Alex Coqueiro, *High Performance Computing in AWS – Technical Workshop*, AMAZON WEST SERVICES PRESENTATION (2015); *Getting Maximum Performance from Amazon Redshift: High Concurrency, Fast Data Loads, and Complex Queries*, AWS RE: INVENT PRESENTATION (November 13, 2013); *Amazon Redshift Database Developer Guide API Version*, AWS DEVELOPER DOCUMENTATION (December 1, 2012); Tony Gibbs, *Amazon Redshift Deep Dive – Storage Subsystem and Query Life Cycle*, AWS PRESENTATION (February 2017); *Amazon Redshift Masterclass – Amazon Web Services*, AWS WEBINAR CHANNEL (October 28, 2015), available at: <https://www.youtube.com/watch?v=GgLKodmL5xE>; Maor Kleider, *Getting Started with Amazon Redshift*, AWS PRESENTATION (February 2017); and Ran Tessler, *Amazon Redshift Deep Dive*, AWS Presentation (July 23, 2015).

106. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '733 patent.

107. As a result of Amazon's infringement of the '733 patent, Plaintiffs have suffered monetary damages, and seeks recovery in an amount adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty for the use made of the invention by Amazon together with interest and costs as fixed by the Court.

COUNT III
INFRINGEMENT OF U.S. PATENT NO. 6,741,983

108. Plaintiffs reference and incorporate by reference the preceding paragraphs of this Complaint as if fully set forth herein.

109. Amazon designs, makes, uses, sells, and/or offers for sale in the United States products and/or services for processing data within a database whereby tests based upon selected vectors are performed on each data record and the data records organized into clusters.

110. Amazon designs, makes, sells, offers to sell, imports, and/or uses Amazon Relational Database Service (RDS) including Amazon RDS for Oracle Versions 11g and 12c and Oracle on Amazon EC2 Versions 11g and 12c (collectively, the "Amazon '983 Product(s)").

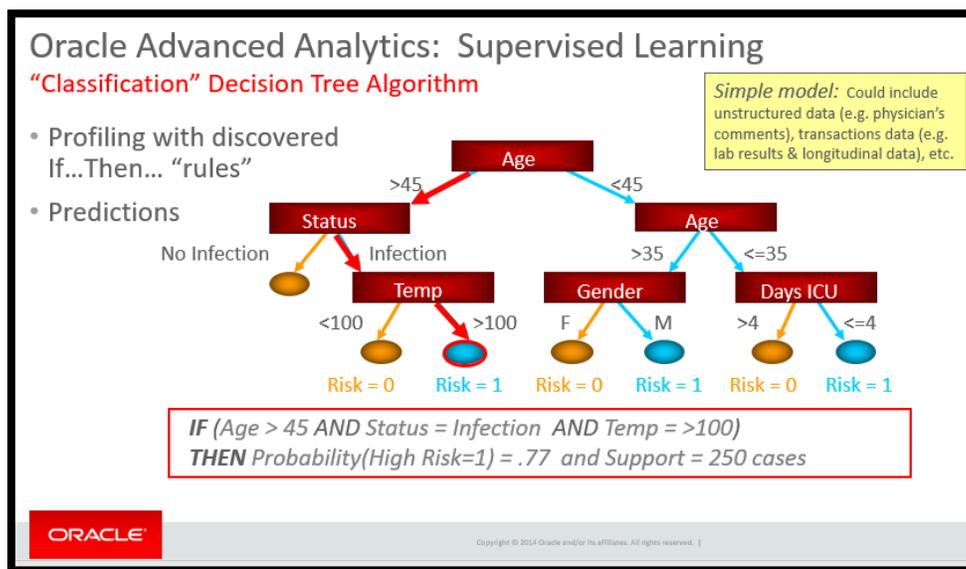
111. On information and belief, one or more Amazon subsidiaries and/or affiliates use the Amazon '983 Products in regular business operations.

112. On information and belief, one or more of the Amazon Products enable organizing data records into clusters.

113. On information and belief, one or more of the Amazon '983 Products include technology for representing one or more variables in each data record in a binary form, whereby the value of each bit is assigned based on the value of a variable.

114. On information and belief, one or more of the Amazon ‘983 Products allow a set of variables to be chosen from those represented in all the data records, whereby principal analysis of the set of variables yields distinct clusters of the data records.

115. On information and belief, one or more of the Amazon ‘983 Products apply principal component analysis to the chosen set of variables for a sample of the data records, whereby two or more principal component vectors are identified.



Charlie Berger, *Big Data Analytics with Oracle Advanced Analytics 12 and Big Data SQL*, Oracle, ORACLE WEBCAST PRESENTATION at 8 (2014).

116. On information and belief, one or more of the Amazon ‘983 Products select vectors of the two more vectors for which the examined scores form distinct clusters.

117. On information and belief, one or more of the Amazon ‘983 Products formulate a test based on selected vectors.

118. On information and belief, one or more of the Amazon ‘983 Products perform a test based upon selected vectors on each data record, whereby the data records are organized into clusters.

119. On information and belief, the Amazon '983 Products are available to businesses and individuals throughout the United States.

120. On information and belief, the Amazon '983 Products are provided to businesses and individuals located in the Eastern District of Tennessee.

121. On information and belief, Amazon has directly infringed and continues to directly infringe the '983 patent by, among other things, making, using, offering for sale, and/or selling technology for extracting data from sources of network-based information, including but not limited to the Amazon '983 Products, which include infringing technology for processing data within a database whereby tests based upon selected vectors are performed on each data record and the data records organized into clusters. Such products and/or services include, by way of example and without limitation, the Amazon '983 Products.

122. By making, using, testing, offering for sale, and/or selling products and services, including but not limited to the Amazon '983 Products, Amazon has injured Plaintiffs and is liable for directly infringing one or more claims of the '983 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

123. On information and belief, Amazon also indirectly infringes the '983 patent by actively inducing infringement under 35 USC § 271(b).

124. On information and belief, Amazon has had knowledge of the '983 patent since at least September 17, 2014. On September 17, 2014, the University of Tennessee Research Foundation contacted Amazon via email in an attempt to secure a license to the patents-in-suit outside the context of litigation.

125. Alternatively, Amazon has had knowledge of the '983 patent since at least service of this Complaint or shortly thereafter, and on information and belief, Amazon knew of the '983 patent and knew of its infringement, including by way of this lawsuit.

126. On information and belief, Amazon intended to induce patent infringement by third-party customers and users of the Amazon '983 Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Amazon specifically intended and was aware that the normal and customary use of the accused products would infringe the '983 patent. Amazon performed the acts that constitute induced infringement, and would induce actual infringement, with knowledge of the '983 patent and with the knowledge that the induced acts would constitute infringement. For example, Amazon provides the Amazon '983 Products that have the capability of operating in a manner that infringe one or more of the claims of the '983 patent, including at least claim 1, and Amazon further provides documentation and training materials that cause customers and end users of the Amazon '983 Products to utilize the products in a manner that directly infringe one or more claims of the '983 patent.³² By providing instruction and training to customers and end-users on how to use the Amazon '983 Products in a manner that directly infringes one or more claims of the '983 patent, including at least claim 1, Amazon specifically intended to induce infringement of the '983 patent. On information and belief, Amazon engaged in such inducement to promote the sales of the Amazon '983 Products, e.g.,

³² See e.g., *Amazon Relational Database Service User Guide- API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); *Overview of AWS Security – Database Services*, AWS WHITE PAPER (June 2016); *Amazon Relational Database API Reference – API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); and *Oracle Advanced Architectures for Oracle Database on Amazon EC2*, ORACLE WHITE PAPER (April 2016).

through Amazon user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '983 patent. Accordingly, Amazon has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '983 patent, knowing that such use constitutes infringement of the '983 patent.

127. The '983 patent is well-known within the industry as demonstrated by the over 85 citations to the '983 patent in published patents and patent applications assigned to technology companies and academic institutions. Amazon is utilizing the technology claimed in the '983 patent without paying a reasonable royalty. Amazon is infringing the '983 patent in a manner best described as willful, wanton, malicious, in bad faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate.

128. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '983 patent.

129. As a result of Amazon's infringement of the '983 patent, Plaintiffs have suffered monetary damages, and seek recovery in an amount adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty for the use made of the invention by Amazon together with interest and costs as fixed by the Court.

COUNT IV
INFRINGEMENT OF U.S. PATENT NO. 7,272,612

130. Plaintiffs reference and incorporate by reference the preceding paragraphs of this Complaint as if fully set forth herein.

131. Amazon designs, makes, uses, sells, and/or offers for sale in the United States products and/or services for partitioning data within a database system by using a function of a

distribution of values that comprise a combination of measures of entropy and adjacency, adjacency being weighted by a weighting factor.

132. Amazon designs, makes, sells, offers to sell, imports, and/or uses Amazon Machine Learning with Amazon Redshift; Amazon Relational Database Service (RDS), including Amazon RDS for Oracle Versions 11g and 12c; and Oracle on Amazon EC2 Versions 11g and 12c (collectively, the “Amazon ‘612 Product(s)’”).

133. On information and belief, one or more Amazon subsidiaries and/or affiliates use the Amazon ‘612 Products in regular business operations.

134. On information and belief, one or more of the Amazon ‘612 Products include technology for partitioning data within a database by defining a function of a distribution of values of a designated variable associated with the data records. Specifically, the Amazon ‘612 Products calculate an entropy measure for a data record. Further, the Amazon ‘612 Products perform this function wherein the function comprises a combination of measures of entropy and adjacency, adjacency being weighted by a weighting factor. For example, Amazon documentation states:

The learning algorithm’s task is to learn the weights for the model. The weights describe the likelihood that the patterns that the model is learning reflect actual relationships in the data. A learning algorithm consists of a loss function and an optimization technique. The loss is the penalty that is incurred when the estimate of the target provided by the ML model does not equal the target exactly. A loss function quantifies this penalty as a single value. An optimization technique seeks to minimize the loss. In Amazon Machine Learning, we use three loss functions, one for each of the three types of prediction problems. The optimization technique used in Amazon ML is online Stochastic Gradient Descent (SGD). SGD makes sequential passes over the training data, and during each pass, updates feature weights one example at a time with the aim of approaching the optimal weights that minimize the loss.

Learning Algorithm, AWS Documentation, available at <http://docs.aws.amazon.com/machine-learning/latest/dg/learning-algorithm.html> (last visited April 2017) (emphasis added).

135. On information and belief, one or more of the Amazon ‘612 Products include technology for partitioning the values of the designated variable into two or more groups, wherein a value of the function is determined by applying an optimization procedure.

Term	Definition
Regression	The goal of training a regression ML model is to predict a numeric value.
Multiclass	The goal of training a multiclass ML model is to predict values that belong to a limited, pre-defined set of permissible values.
Binary	The goal of training a binary ML model is to predict values that can only have one of two states, such as true or false.
Model Size	ML models capture and store patterns. The more patterns a ML model stores, the bigger it will be. ML model size is described in Mbytes.
Number of Passes	When you train an ML model, you use data from a datasource. It is sometimes beneficial to use each data record in the learning process more than once. The number of times that you let Amazon ML use the same data records is called the number of passes.
Regularization	Regularization is a machine learning technique that you can use to obtain higher-quality models. Amazon ML offers a default setting that works well for most cases.

Amazon Machine Learning Key Concepts, AMAZON MACHINE LEARNING DEVELOPER GUIDE at 3 (2017) (describing how Amazon ML supports three types of ML models: binary classification, multiclassification and regression).

136. On information and belief, one or more of the Amazon ‘612 Products include technology for assigning a data record to a group according to the values of the designated variable.

137. On information and belief, the Amazon ‘612 Products are available to businesses and individuals throughout the United States.

138. On information and belief, the Amazon ‘612 Products are provided to businesses and individuals located in the Eastern District of Tennessee.

139. On information and belief, Amazon has directly infringed and continues to directly infringe the ‘612 patent by, among other things, making, using, offering for sale, and/or selling technology for extracting data from sources of network-based information, including but not limited to the Amazon ‘612 Products, which include infringing technology for partitioning data within a database system by using a function of a distribution of values that comprise a

combination of measures of entropy and adjacency, adjacency being weighted by a weighting factor. Such products and/or services include, by way of example and without limitation, the Amazon '612 Products.

140. By making, using, testing, offering for sale, and/or selling products and services, including but not limited to the Amazon '612 Products, Amazon has injured Plaintiffs and is liable for directly infringing one or more claims of the '612 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

141. On information and belief, Amazon also indirectly infringes the '612 patent by actively inducing infringement under 35 U.S.C. § 271(b).

142. On information and belief, Amazon has had knowledge of the '612 patent since at least September 17, 2014. On September 17, 2014, the University of Tennessee Research Foundation contacted Amazon via email in an attempt to secure a license to the patents-in-suit outside the context of litigation.

143. Alternatively, on information and belief Amazon has had knowledge of the '612 patent since at least service of this Complaint or shortly thereafter, and on information and belief, Amazon knew of the '612 patent and knew of its infringement, including by way of this lawsuit.

144. On information and belief, Amazon intended to induce patent infringement by third-party customers and users of the Amazon '612 Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Amazon specifically intended and was aware that the normal and customary use of the accused products would infringe the '612 patent. Amazon performed the acts that constitute induced infringement, and would induce actual infringement, with knowledge of the '612 patent and with the knowledge that the induced acts would constitute

infringement. For example, Amazon provides the Amazon ‘612 Products that have the capability of operating in a manner that infringe one or more of the claims of the ‘612 patent, including at least claim 1, and Amazon further provides documentation and training materials that cause customers and end users of the Amazon ‘612 Products to utilize the products in a manner that directly infringe one or more claims of the ‘612 patent.³³ By providing instruction and training to customers and end-users on how to use the Amazon ‘612 Products in a manner that directly infringes one or more claims of the ‘612 patent, including at least claim 1, Amazon specifically intended to induce infringement of the ‘612 patent. On information and belief, Amazon engaged in such inducement to promote the sales of the Amazon ‘612 Products, e.g., through Amazon user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the ‘612 patent. Accordingly, Amazon has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ‘612 patent, knowing that such use constitutes infringement of the ‘612 patent.

145. The ‘612 patent is well-known within the industry as demonstrated by the over 26 citations to the ‘612 patent in published patents and patent applications assigned to technology companies and academic institutions. Amazon is utilizing the technology claimed in the ‘612 patent without paying a reasonable royalty. Amazon is infringing the ‘612 patent in a manner

³³ See e.g., *Amazon Machine Learning Key Concepts*, AMAZON MACHINE LEARNING DEVELOPER GUIDE (2017); John Chang, *Building a Recommendation Engine and Use Amazon ML in Real Time*, AWS PRESENTATION (September 2016); *Amazon Relational Database Service User Guide- API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); *Overview of AWS Security – Database Services*, AWS WHITE PAPER (June 2016); *Amazon Relational Database API Reference – API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); and *Oracle Advanced Architectures for Oracle Database on Amazon EC2*, ORACLE WHITE PAPER (April 2016).

best described as willful, wanton, malicious, in bad faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate.

146. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '612 patent.

147. As a result of Amazon's infringement of the '612 patent, Plaintiffs have suffered monetary damages, and seek recovery in an amount adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty for the use made of the invention by Amazon together with interest and costs as fixed by the Court.

COUNT V
INFRINGEMENT OF U.S. PATENT NO. 7,882,106

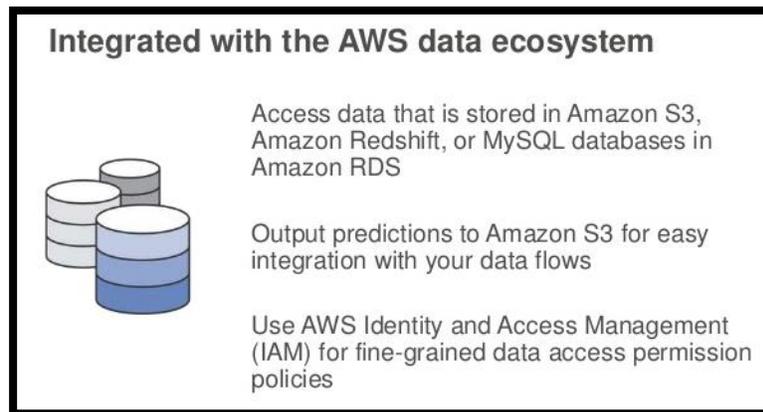
148. Plaintiffs reference and incorporate by reference the preceding paragraphs of this Complaint as if fully set forth herein.

149. Amazon designs, makes, uses, sells, and/or offers for sale in the United States products and/or services for partitioning data records of multi-dimensional databases utilizing a parallel data processing system wherein the records are partitioned using a function that comprises a combination of measures of entropy and adjacency, adjacency being weighted by a weighting factor.

150. Amazon designs, makes, sells, offers to sell, imports, and/or uses Amazon Machine Learning with Amazon Redshift; Amazon Relational Database Service (RDS) including Amazon RDS for Oracle Versions 11g and 12c; and Oracle on Amazon EC2 Versions 11g and 12c (collectively, the "Amazon '106 Product(s)").

151. On information and belief, one or more Amazon subsidiaries and/or affiliates use the Amazon '106 Products in regular business operations.

152. On information and belief, the Amazon ‘106 Products comprise a parallel data processing system with first and second computer processors for implementing a method of partitioning data records of a multi-dimensional database into groups.



John Chang, *Building a Recommendation Engine and Use Amazon ML in Real Time*, AWS PRESENTATION at 28 (September 2016).

153. On information and belief, one or more of the Amazon ‘106 Products include technology for defining a function of a distribution of values of a designated variable associated with the data records, wherein the function comprises a combination of measures of entropy and adjacency where adjacency is weighted by a weighting factor.

The learning algorithm’s task is to learn the weights for the model. The weights describe the likelihood that the patterns that the model is learning reflect actual relationships in the data. A learning algorithm consists of a loss function and an optimization technique. The loss is the penalty that is incurred when the estimate of the target provided by the ML model does not equal the target exactly. A loss function quantifies this penalty as a single value. An optimization technique seeks to minimize the loss. In Amazon Machine Learning, we use three loss functions, one for each of the three types of prediction problems. The optimization technique used in Amazon ML is online Stochastic Gradient Descent (SGD). SGD makes sequential passes over the training data, and during each pass, updates feature weights one example at a time with the aim of approaching the optimal weights that minimize the loss.

Learning Algorithm, AWS Documentation, available at <http://docs.aws.amazon.com/machine-learning/latest/dg/learning-algorithm.html> (last visited April 2017) (emphasis added).

154. On information and belief, one or more of the Amazon ‘106 Products include technology for partitioning the values of the designated variable into two or more groups for storage in computer processor memory.

155. On information and belief, the Amazon ‘106 Products are available to businesses and individuals throughout the United States.

156. On information and belief, the Amazon ‘106 Products are provided to businesses and individuals located in the Eastern District of Tennessee.

157. On information and belief, Amazon has directly infringed and continues to directly infringe the ‘106 patent by, among other things, making, using, offering for sale, and/or selling technology for extracting data from sources of network-based information, including but not limited to the Amazon ‘106 Products, which include infringing technology for partitioning data records of multi-dimensional databases utilizing a parallel data processing system wherein the records are partitioned using a function that comprises a combination of measures of entropy and adjacency, adjacency being weighted by a weighting factor. Such products and/or services include, by way of example and without limitation, the Amazon ‘106 Products.

158. By making, using, testing, offering for sale, and/or selling products and services, including but not limited to the Amazon ‘106 Products, Amazon has injured Plaintiffs and is liable for directly infringing one or more claims of the ‘106 patent, including at least claim 14, pursuant to 35 U.S.C. § 271(a).

159. On information and belief, Amazon also indirectly infringes the ‘106 patent by actively inducing infringement under 35 USC § 271(b).

160. Alternatively, on information and belief, Amazon has had knowledge of the ‘106 patent since at least September 17, 2014. On September 17, 2014, the University of Tennessee

Research Foundation contacted Amazon via email in an attempt to secure a license to the patents-in-suit outside the context of litigation.

161. On information and belief, Amazon has had knowledge of the ‘106 patent since at least service of this Complaint or shortly thereafter, and on information and belief, Amazon knew of the ‘106 patent and knew of its infringement, including by way of this lawsuit.

162. On information and belief, Amazon intended to induce patent infringement by third-party customers and users of the Amazon ‘106 Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Amazon specifically intended and was aware that the normal and customary use of the accused products would infringe the ‘106 patent. Amazon performed the acts that constitute induced infringement, and would induce actual infringement, with knowledge of the ‘106 patent and with the knowledge that the induced acts would constitute infringement. For example, Amazon provides the Amazon ‘106 Products that have the capability of operating in a manner that infringe one or more of the claims of the ‘106 patent, including at least claim 14, and Amazon further provides documentation and training materials that cause customers and end users of the Amazon ‘106 Products to utilize the products in a manner that directly infringe one or more claims of the ‘106 patent.³⁴ By providing instruction and training to customers and end-users on how to use the Amazon ‘106 Products in a manner

³⁴ See, e.g., *Amazon Machine Learning Key Concepts*, AMAZON MACHINE LEARNING DEVELOPER GUIDE (2017); John Chang, *Building a Recommendation Engine and Use Amazon ML in Real Time*, AWS PRESENTATION (September 2016); *Amazon Relational Database Service User Guide- API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); *Overview of AWS Security – Database Services*, AWS WHITE PAPER (June 2016); *Amazon Relational Database API Reference – API Version 2014-10-31*, AMAZON DEVELOPER DOCUMENTATION (October 31, 2014); and *Oracle Advanced Architectures for Oracle Database on Amazon EC2*, ORACLE WHITE PAPER (April 2016).

that directly infringes one or more claims of the '106 patent, including at least claim 14, Amazon specifically intended to induce infringement of the '106 patent. On information and belief, Amazon engaged in such inducement to promote the sales of the Amazon '106 Products, e.g., through Amazon user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '106 patent. Accordingly, Amazon has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '106 patent, knowing that such use constitutes infringement of the '106 patent.

163. The '106 patent is well-known within the industry as demonstrated by the multiple citations to the '106 patent in published patents and patent applications assigned to technology companies and academic institutions. Amazon is utilizing the technology claimed in the '106 patent without paying a reasonable royalty. Amazon is infringing the '106 patent in a manner best described as willful, wanton, malicious, in bad faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate.

164. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '106 patent.

165. As a result of Amazon's infringement of the '106 patent, Plaintiffs have suffered monetary damages, and seek recovery in an amount adequate to compensate for Amazon's infringement, but in no event less than a reasonable royalty for the use made of the invention by Amazon together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs respectfully requests that this Court enter:

- A. A judgment in favor of Plaintiffs that Amazon has infringed, either literally and/or under the doctrine of equivalents, the ‘411, ‘733, ‘983, ‘612, and ‘106 patents;
- B. An award of damages resulting from Amazon’s acts of infringement in accordance with 35 U.S.C. § 284;
- C. A judgment and order finding that Defendant’s infringement was willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate within the meaning of 35 U.S.C. § 284 and awarding to Plaintiffs enhanced damages.
- D. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiffs their reasonable attorneys’ fees against Defendant.
- E. Any and all other relief to which Plaintiffs may show themselves to be entitled.

JURY TRIAL DEMANDED

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, University of Tennessee Research Foundation and Saint Matthew Research, LLC requests a trial by jury of any issues so triable by right.

Dated: May 2, 2017

Respectfully submitted,

s/Wayne A. Ritchie II
WAYNE A. RITCHIE II (BPR 013936)
JAMES R. STOVALL (BPR #032512)
RITCHIE, DILLARD, DAVIES
& JOHNSON, P.C.
606 West Main Avenue, Suite 300
P.O. Box 1126
Knoxville, Tennessee 37901-1126
(865) 637-0661
E-mail: war@rddjlaw.com
E-mail: jstovall@rddjlaw.com

OF COUNSEL:

Dorian S. Berger (CA SB No. 264424)
pro hac vice application forthcoming
Daniel P. Hipskind (CA SB No. 266763)
pro hac vice application forthcoming
Eric B. Hanson (CA SB No. 254570)
pro hac vice application forthcoming
BERGER & HIPSKIND LLP
1880 Century Park East, Ste. 815
Los Angeles, CA 95047
Telephone: 323-886-3430
Facsimile: 323-978-5508
E-mail: dsb@bergerhipskind.com
E-mail: dph@bergerhipskind.com
E-mail: ebh@bergerhipskind.com

*Attorneys for Plaintiffs University of
Tennessee Research Foundation and Saint
Matthew Research, LLC*