

**IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF TEXAS
FORT WORTH DIVISION**

VALTRUS INNOVATIONS LTD.,

Plaintiff,

v.

GOOGLE LLC,

Defendant.

Case No. 4:22-cv-00020

JURY TRIAL DEMANDED

VALTRUS INNOVATIONS LTD.'S COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Valtrus Innovations Limited (“Valtrus”), by and through its undersigned counsel, pleads the following against Google LLC (“Google”) and alleges as follows:

THE PARTIES

1. Plaintiff Valtrus is the successor in interest to a substantial patent portfolio created by Hewlett Packard Enterprise and its predecessor companies, including Compaq, Verity, and Hewlett-Packard Development Company (collectively, “HPE”). Valtrus is an Irish entity duly organized and existing under the laws of the Republic of Ireland. The address of the registered office of Valtrus is: The Glasshouses GH2, 92 Georges Street Lower, Dun Laoghaire, Dublin A96 VR66, Ireland. HPE’s worldwide corporate headquarters is located in Houston, Texas. One of HPE’s primary US facilities is located in Plano, Texas.

2. Valtrus is the assignee and owns all right and title to U.S. Patent Nos. 6,728,704 (“the ’704 Patent”), 6,738,764 (“the ’764 Patent”), 6,816,809 (“the ’809 Patent”), 7,346,604 (“the ’604 Patent”), 7,523,454 (“the ’454 Patent”), and 7,748,005 (“the ’005 Patent”) (collectively, the “Asserted Patents”).

3. The Asserted Patents were developed by inventors working for HPE. HPE and its predecessors have been developing innovative search, computer processing, and server technology for decades.

4. On information and belief, Defendant Google is a limited liability company duly organized and existing under the laws of the State of Delaware, having a regular and established place of business in the Northern District of Texas, including at 3800 Railport Parkway, Midlothian, Texas 76065.

5. Google has become the dominant search engine in the United States and beyond, with a reported market share of around 90%. Google’s strategy has included taking technology from other companies without regard to intellectual property rights and combining that

technology to create and maintain market dominance. A core part of this strategy involved taking the innovations in the Asserted Patents. One need look no further than Google's employment records, which have listed as some of Google's most senior technologists a large number of inventors on the Asserted Patents, including a Senior Vice President at Google responsible for search, the first Director of Google R&D Bangalore, and the Vice President at Google responsible for Google Apps and Cloud.

JURISDICTION AND VENUE

6. This is an action arising under the patent laws of the United States, 35 U.S.C. § 1 *et seq.* Accordingly, this Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

7. This Court has personal jurisdiction over Google because Google creates products and services that are and have been used, offered for sale, sold, and purchased in the Northern District of Texas, and Google has committed, and continues to commit, acts of infringement in the Northern District of Texas, has conducted business in the Northern District of Texas, and/or has engaged in continuous and systematic activities in the Northern District of Texas.

8. Under 28 U.S.C. §§ 1391(b)-(d) and 1400(b), venue is proper in this judicial district because Google maintains a regular and established place of business in this district and has committed and regularly commits acts of infringement within this judicial district giving rise to this action. For example, Google operates a 260,000 square foot data center in Midlothian, Texas. Valtrus is informed and believes that this data center includes the infringing systems and practices the infringing methods described herein. This data center is one of only fourteen in all of North America. Valtrus is informed and believes that Google is preparing to open a Google Cloud location in Dallas, which will be one of only eleven Cloud locations in North America, and which similarly is and/or will be a location of significant infringement. On information and

belief, Google also maintains an office in the Dallas-Fort Worth area, and runs a content distribution network (CDN) node in the Dallas area. Google also provides, sells, and offers for sale infringing products and services to users in the Northern District of Texas.

9. Google also has significant operations in nearby cities, including Austin. For example, on information and belief, Google owns approximately 550,000 square feet of office space across three locations in downtown Austin, Texas. Google is also continuing to grow its presence in Texas. For example, Google is preparing to open an additional 750,000 square feet of offices in Austin at Block 185, a new office tower located at 601 West Second Street that is set to open in 2023. Valtrus is informed and believes that these offices include employees responsible for the subject matter of this patent suit. For example, on information and belief, Google employees in Austin work on Google Cloud, finance, and engineering. On information and belief, Google also has an office in Houston, and employs more than 1700 people in Texas.

FIRST CLAIM

(Infringement of U.S. Patent No. 6,728,704)

10. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-9 of its Complaint.

11. The '704 Patent, entitled "Method and apparatus for merging result lists from multiple search engines," was duly and lawfully issued on April 27, 2004. A true and correct copy of the '704 Patent is attached hereto as Exhibit 1.

12. The '704 Patent names Jianchang Mao, Rajat Mukherjee, Prabhakar Raghavan, and Panayiotis Tsaparas as co-inventors.

13. The '704 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '704 Patent, including the right to seek damages for any infringement thereof.

14. The '704 Patent “relates generally to search engine technology. More specifically, this invention relates to reducing the computational overhead associated with merging results from multiple search engines.” Ex. 1 at 1:7-10.

15. The '704 Patent explains that “the merging of multiple result lists into a single list is usually accomplished by examining and ranking every single entry of every list. . . . Thus, for large lists or large numbers of lists, the computation time required by the merging process can nullify any advantage gained by operating multiple search engines at the same time.” *Id.* at 2:48-51, 53-56.

16. The '704 Patent states that the “invention allows for a reduction in computational overhead when merging and re-ranking multiple result lists. Ranking of results is accomplished by evaluating a subset of entries instead of every single one, thus reducing the number of calculations required.” *Id.* at 3:20-24.

17. The method described by the '704 Patent improves the function of a computer utilizing said method by reducing the considerable computational overhead associated with merging many results from numerous search engines. A reduction in computational overhead enables the computer to operate multiple search engines simultaneously without sacrificing response time, allowing the gathering, ranking, and presentation of millions of search results from across the Internet in a very short amount of time. The '704 Patent provides detailed, specific steps for merging multiple result lists in this manner. For example, in one embodiment, each result list is assigned a representative scoring value based on a selected subset of entries from the list. *Id.* at 5:44-55; 7:12-14. In another embodiment, each list is instead “assigned a probability value equal to its average scoring value’s percentage of the total of all average


scoring values,” and entries are selected from each list based on its probability value. *Id.* at 7:35-39.

18. Valtrus is informed and believes, and thereon alleges, that Google has infringed and unless enjoined will continue to infringe one or more claims of the ’704 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed method of merging result lists in an infringing manner. Google practices every step of at least claim 1 of the ’704 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

19. For example, the ’704 accused product, Google Search, embodies every limitation of at least claim 1 of the ’704 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.


20. Google Search practices a method of merging result lists from multiple search engines, comprising the elements set forth below.

21. For example, a query transmitted to Google Search returns results from multiple search engines displayed to the user as a single merged list of results. *See, e.g.,* Google Search Results for “in-n-out burger” Query 1 (red arrows added):




22. Google Search practices a method comprising transmitting a query to a set of search engines.

23. For example, a query transmitted to Google Search returns results from multiple search engines. This process begins by transmitting that query to said search engines. *See, e.g., id.* (red arrows added):




24. Google Search practices a method comprising receiving in response to said query a result list from each search engine of said set of search engines, each result list including one or more entries.

25. For example, a query transmitted to Google Search returns one or more results from each of several search engines. Said results are grouped into result lists associated with each of said search engines. In order to display said result lists, Google Search first receives a result list from each search engine. *See, e.g., id.* (red arrows added):




26. Google Search practices a method comprising selecting a subset of entries from each result list to form a set of selected entries.

27. For example, a query transmitted to Google Search may return many millions of results across a set of multiple search engines. Only a subset of these results, or entries, is selected for display to the user near the top of the Google Search result list. *See, e.g.*, Google Search Results for “in-n-out burger” Query 2 (callout added):




28. Google Search practices a method comprising assigning to each selected entry of said set of selected entries a scoring value according to a scoring function.

29. Entries are displayed to a user of Google Search in order of relevance. On information and belief, these entries are assigned a scoring value according to a scoring function in order to be sorted and displayed in this way. For example, a list of results from the Maps subset of entries is sorted according to a scoring function which considers metrics such as the distance of each result from the user at the time of the query. *See, e.g., Google Maps Result List* (red arrow added):




30. Google Search practices a method comprising assigning to each subset a representative value according to the scoring values assigned to its entries.

31. For example, Google Search presents results in subsets that appear in different orders for different queries. On information and belief, Google Search orders each subset based on a representative value according to the scoring values assigned to its entries. *See, e.g., Google Search Results for “in-n-out burger” Query 1* (red arrows added):



32. Google Search practices a method comprising producing a merged list of entries in a predetermined manner based on the representative value assigned to each result list.


33. For example, when a user transmits a query involving a term associated with a physical location, Google Search produces a merged list of entries where in Maps results are generally displayed higher in the merged list than News results because, on information and belief, the Maps result list has been assigned a higher representative value. *See, e.g., id.* (red arrows added):



34. Google Search practices a method comprising producing a merged list of entries in a predetermined manner based on the representative value assigned to each result list, wherein the representative value varies in accordance with predetermined manner.

35. For example, result lists from each individual search engine are generally displayed as individual groups of results in the merged list presented to the user. On information and belief, the representative value of each result list, which determines the order of said groups,

varies in accordance with the nature of a user's query. *See, e.g.,* Google Search Result Groupings (red arrows added):



36. Google has had knowledge of the '704 Patent since at least January 1, 2013. On this date, two patents assigned to Google Inc. were issued, both of which cite the '704 Patent as a reference: U.S. Patent Nos. 8,346,791 and 8,346,792. The '704 Patent is also cited as a reference by two later patents, U.S. Patent Nos. 8,521,725 (issued August 27, 2013) and 9,152,714 (issued October 6, 2015), both of which name Google Inc. as assignee. Furthermore, two of the co-inventors of the '704 Patent, Rajat Mukherjee and Prabhakar Raghavan, are listed as inventors on at least seven other patents or publications which are assigned to Google: U.S. Patent Nos.

8,782,071 (issued July 15, 2014), 9,251,168 (issued February 2, 2016), 10,079,785 (issued September 18, 2018), 10,402,889 (issued September 3, 2019), and 10,917,371 (issued February 9, 2021); and U.S. Patent Application Publication Nos. 2014/0258466 (published September 11, 2014) and 2016/0371425 (published December 22, 2016). Two of these patents, U.S. Patent Nos. 8,782,071 and 10,402,889, relate directly to search technology. Furthermore, Prabhakar Raghavan now holds the position of Senior Vice President at Google, where he “is responsible for Google Search,” among other products.¹ Given that four of its own patents cite the ’704 Patent as a reference, two of the co-inventors of the ’704 Patent are listed as inventors on numerous later patents assigned to Google, and one of the co-inventors of the ’704 Patent is now in a senior leadership role at Google, Google knew or was willfully blind to the fact that its conduct was infringing by, at the very latest, July 15, 2014.

37. Valtrus is informed and believes, and thereon alleges, that Google actively, knowingly, and intentionally has induced infringement of the ’704 Patent by, for example, offering for public use Google Search with the intent to encourage and facilitate infringing uses of that service in the Northern District of Texas, in the United States, and throughout the world.

38. As a result of Google’s infringement of the ’704 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google’s wrongful acts in an amount subject to proof at trial.

39. In addition, Google’s infringing acts and practices have caused and are causing immediate and irreparable harm to Valtrus.

40. Valtrus is informed and believes, and thereon alleges, that Google’s infringement of the ’704 Patent has been and continues to be willful. As noted above, Google has long had

¹ <https://research.google/people/PrabhakarRaghavan/>

knowledge of the '704 Patent and its infringement of the '704 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Valtrus's patent rights. Thus, Google's infringing actions have been and continue to be consciously wrongful.

SECOND CLAIM

(Infringement of U.S. Patent No. 6,738,764)

41. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-40 of its Complaint.

42. The '764 Patent, entitled "Apparatus and method for adaptively ranking search results," was duly and lawfully issued on May 18, 2004. A true and correct copy of the '764 Patent is attached hereto as Exhibit 2.

43. The '764 Patent names Jianchang Mao, Mani Abrol, Rajat Mukherjee, Michel Tourn, and Prabhakar Raghavan as co-inventors.

44. The '764 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '764 Patent, including the right to seek damages for any infringement thereof.

45. The '764 Patent "relates generally to computerized techniques for identifying relevant documents. More particularly, this invention describes computerized techniques for adaptively ranking documents identified in response to a search query." Ex. 2 at 1:6-10.

46. The '764 Patent states that "it is impossible to predict a priori, for any corpus of documents and any associated search engine, on which queries the static method [of ranking search results] is satisfactory and on which other queries the adaptive method is satisfactory." *Id.* at 1:51-55.

47. The '764 Patent explains that it is “highly desirable to provide a technique that selectively emphasizes a static method or an adaptive method to achieve optimal search results for a given query.” *Id.* at 1:56-59. “The invention provides improved search results by adaptively ranking, based upon the prior behavior of users, documents returned from a text search engine. More particularly, the prior behavior of users is utilized to determine the rate at which to apply adaptive correction for a given query.” *Id.* at 2:30-35.

48. The method described by the '764 Patent provides a novel approach for selectively emphasizing a static or an adaptive method of ranking search results based, at least in part, on the prior behavior of other users. By adaptively ranking documents based, at least in part, on other users' behavior and building “a list of queries associated with a viewed document,” *id.* at 4:30, this method improves on the performance and operation of then-existing computer systems, which were unable “to predict a priori, for any corpus of documents and any associated search engine, on which queries the static method is satisfactory and on which other queries the adaptive method is satisfactory.” *Id.* at 1:52-55. The '764 Patent describes specific steps for accomplishing this improvement through a number of different implementations. For example, the '764 Patent sets out in detail a series of processing operations associated with an embodiment of the invention. *Id.* at 4:2-3. These operations include producing “relevance search results based upon a query,” logging a user's pattern of viewing search result documents, producing a database “which stores a list of queries associated with a viewed document,” forming a feature vector for each viewed document which “characterizes attributes and query words associated” with it, and calculating a similarity score for a user's query using said feature vectors. *Id.* at 4:4-46. These and other detailed descriptions provide step-by-step instructions for carrying out the method and thereby improve the function of a computer using the method. This method


significantly reduces computational overhead and improves the speed of return of the search results.

49. Valtrus is informed and believes, and thereon alleges, that Google has infringed and unless enjoined will continue to infringe one or more claims of the '764 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed method of adaptively ranking search results in an infringing manner. Google practices every step of at least claim 1 of the '764 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

50. For example, the '764 accused product, Google Search, embodies every limitation of at least claim 1 of the '764 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

51. Google Search practices a method of ranking search results comprising the elements described below.

52. When a user transmits a query, Google Search ranks search results. *See, e.g.*, How Search works – Ranking results² (red rectangular annotation added):



With the vast amount of information available, finding what you need would be nearly impossible without some help sorting through it. Google's ranking systems are designed to do just that; sort through hundreds of billions of webpages and other content in our Search index to present the most relevant, useful results in a fraction of a second.

² <https://www.google.com/search/howsearchworks/how-search-works/ranking-results/>

53. Google Search practices a method comprising producing a relevance score for a document in view of a query.

54. For example, Google Search algorithms “look at many factors,” including the “relevance and usability of pages,” to determine how to rank search results. In order for said algorithms to rank results, Google Search assigns a relevance score to results. *See, e.g., Google Search consumer information*³ (red underline added):


These ranking systems are made up of not one, but a whole series of algorithms. To give you the most useful information, Search algorithms look at many factors, including the words of your query, relevance and usability of pages, expertise of sources, and your location and settings. The weight applied to each factor varies depending on the nature of your query—for example, the freshness of the content plays a bigger role in answering queries about current news topics than it does about dictionary definitions.

55. Google Search practices a method comprising calculating a similarity score for said query utilizing a feature vector that characterizes attributes and query words of a different query associated with said document.

56. For example, Google Search uses a “neural network-based technique for natural language processing (NLP) pre-training” known as BERT, which can “consider the full context

³ <https://support.google.com/websearch/answer/7585859?hl=en>

of a word by looking at the words that come before and after it.” See Pandu Nayak, *Understanding searches better than ever before*⁴ (red rectangular annotations added):



57. BERT extracts feature vectors, which can characterize attributes and query words of other queries associated with the same site or document for other users. Feature vectors can then be used to determine the degree of similarity between sites and assign a score accordingly.

58. Google Search practices a method comprising assigning a rank value for said document based upon said relevance score and said similarity score.

59. For example, Google Search assigns a rank value to each document produced by a user’s query based on the relevance and similarity scores produced by algorithms including BERT. See, e.g., Google Search consumer information (red underline added):

These ranking systems are made up of not one, but a whole series of algorithms. To give you the most useful information, Search algorithms look at many factors, including the words of your query, relevance and usability of pages, expertise of sources, and your location and settings. The weight applied to each factor varies depending on the nature of your query—for example, the freshness of the content plays a bigger role in answering queries about current news topics than it does about dictionary definitions.

⁴ <https://blog.google/products/search/search-language-understanding-bert/>

60. Google has had knowledge of the '764 Patent at least since May 22, 2007. On this date, U.S. Patent No. 7,222,127, which cites the '764 Patent as a reference, was issued with Google Inc. as assignee. The '764 Patent is also cited as a reference by the following later patents: U.S. Patent Nos. 7,716,225 (issued May 11, 2010), 7,870,147 (issued January 11, 2011), 8,051,096 (issued November 1, 2011), 8,140,524 (issued March 20, 2012), 8,346,791 (issued January 1, 2013), 8,346,792 (issued January 1, 2013), 8,359,309 (issued January 22, 2013), 8,396,865 (issued March 12, 2013), 8,498,974 (issued July 30, 2013), 8,521,725 (issued August 27, 2013), 8,615,514 (issued December 24, 2013), 8,661,029 (issued February 25, 2014), 8,694,511, 8,694,374 (both issued April 8, 2014), 8,832,083 (issued September 9, 2014), 8,843,536 (issued September 23, 2014), 8,874,555 (issued October 28, 2014), 8,909,655 (issued December 9, 2014), 8,924,379 (issued December 30, 2014), 8,938,463 (issued January 20, 2015), 8,959,093 (issued February 17, 2015), 8,972,391, 8,972,394 (both issued March 3, 2015), 9,002,867 (issued April 7, 2015), 9,009,146 (issued April 14, 2015), 9,092,510 (issued July 28, 2015), 9,110,975 (issued August 18, 2015), 9,183,499 (issued November 10, 2015), 9,223,868 (issued December 29, 2015), and 9,623,119 (issued April 18, 2017), all of which name Google Inc. as assignee. Furthermore, three of the co-inventors of the '764 Patent, Rajat Mukherjee, Prabhakar Raghavan, and Michel Tourn, are listed as inventors on at least eight other patents or publications which are assigned to Google: U.S. Patent Nos. 8,738,612 (issued May 27, 2014), 8,782,071 (issued July 15, 2014), 9,251,168 (issued February 2, 2016), 10,079,785 (issued September 18, 2018), 10,402,889 (issued September 3, 2019), and 10,917,371 (issued February 9, 2021); and U.S. Patent Application Publication Nos. 2014/0258466 (published September 11, 2014) and 2016/0371425 (published December 22, 2016). Three of these patents, U.S. Patent Nos. 8,738,612, 8,782,071, and 10,402,889, relate directly to search technology. Furthermore,

Prabhakar Raghavan now holds the position of Senior Vice President at Google, where he “is responsible for Google Search,” among other products.⁵ With more than 30 patents assigned to Google over the course of a decade citing the ’764 Patent as a reference, three of the co-inventors of the ’764 Patent listed as inventors on numerous later patents assigned to Google, and at least one of the co-inventors of the ’764 Patent in a senior leadership role at Google, Google knew or was willfully blind to the fact that its conduct was infringing since at least the earliest issue date of these patents.

61. Valtrus is informed and believes, and thereon alleges, that Google actively, knowingly, and intentionally has induced infringement by, for example, offering for public use Google Search with the intent to encourage and facilitate infringing uses of that service in the Northern District of Texas, in the United States, and throughout the world.

62. As a result of Google’s infringement of the ’764 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google’s wrongful acts in an amount subject to proof at trial.

63. In addition, Google’s infringing acts and practices have caused and are causing immediate and irreparable harm to Valtrus.

64. Valtrus is informed and believes, and thereon alleges, that Google’s infringement of the ’764 Patent has been and continues to be willful. As noted above, Google has long had knowledge of the ’764 Patent and its infringement of the ’764 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Valtrus’s patent rights. Thus, Google’s infringing actions have been and continue to be consciously wrongful.

⁵ <https://research.google/people/PrabhakarRaghavan/>

THIRD CLAIM

(Infringement of U.S. Patent No. 6,816,809)

65. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-64 of its Complaint.

66. The '809 Patent, entitled "Hardware based utilization metering," was duly and lawfully issued on November 9, 2004. A true and correct copy of the '809 Patent is attached hereto as Exhibit 3.

67. The '809 Patent names Edgar Circenis as inventor.

68. The '809 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '809 Patent, including the right to seek damages for any infringement thereof.

69. The '809 Patent's technical field includes "methods that use central processor metering to determine processor utilization for billing and other purposes." Ex. 3 at 1:6-8.

70. The '809 Patent states that, "[i]n a computer system having hardware that may be partitioned, gathering processor utilization data from a hardware system requires communications between the metering application and all operating systems running within the hardware. The need for communication with different operating systems poses significant challenges because operating systems by their design are separated from other operating systems and do not have visibility to utilization data from other operating systems." *Id.* at 1:18-27.

71. The '809 Patent explains that the invention uses "inputs from [a] system clock and [an] idle indicator" to measure CPU cycles "where the CPU is not in an idle state, but instead is performing a service for the user of the system." *Id.* at 4:1-4.

72. Valtrus is informed and believes, and thereon alleges, that Google has infringed and unless enjoined will continue to infringe at least one or more of the method claims of the

'809 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed hardware based utilization metering method in an infringing manner. Google practices every step of at least claim 13 of the '809 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

73. For example, the '809 accused product, Google Cloud, embodies every limitation of at least claim 13 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

74. Google Cloud practices a hardware based method for measuring processor utilization in a computer system comprising a plurality of processors. The method comprises the elements described below.

75. For example, Google Cloud offerings are powered by a plurality of processors. *See, e.g., Confidential Computing*⁶:

Real-time encryption in use

Google Cloud customers can encrypt data in use, taking advantage of security technology offered by modern CPUs (e.g., Secure Encrypted Virtualization extension supported by 2nd Gen AMD EPYC™ CPUs) together with confidential computing cloud services. Customers can be confident that their data will stay private and encrypted even while being processed.

76. Google Cloud practices a method comprising determining when any of the plurality of processors is busy.

⁶ <https://cloud.google.com/confidential-computing>

77. For example, processors used by Google Cloud have performance counters. An exemplary performance counter in these processors increments only while the processor core is in the C0 state, meaning the core is executing instructions, and is used to determine when any of a plurality of such processors is busy.

78. Google Cloud practices a method comprising providing a busy indication to a counter associated with a busy processor.

79. For example, the processor cores used by Google Cloud provide a busy indication to a counter while the core is in the C0 state.

80. Google Cloud practices a method comprising receiving at the counter a measure of computer system time.

81. For example, a counter in the processors used by Google Cloud receives at the counter a core clock while the core is in the C0 state.

82. Google Cloud practices a method comprising incrementing a counter value in the counter based on the provided busy indication and an amount of computer system time that the processor is determined to be busy.

83. For example, counters in the processors used by Google Cloud are incremented by hardware based on busy indications and the core clock.

84. Google Cloud practices a method comprising maintaining the counter value.

85. For example, counters in the processors used by Google Cloud are read/write registers, meaning that their values are preserved for access by software.

86. Google has had knowledge of the '809 Patent at least since April 14, 2021. On that date, Valtrus sent to Google a notice letter with an exemplary list of Valtrus patents, which included the '809 Patent, that were infringed by specific exemplary Google products. Therefore,

Google either knew of or was willfully blind to the '809 Patent and its infringement of the '809 Patent no later than this date.

87. Valtrus is informed and believes, and thereon alleges, that Google actively, knowingly, and intentionally has induced infringement of the '809 Patent by, for example, using, selling, and offering for sale Google Cloud products which rely on processors having the claimed counters, which in turn use the process claimed by the '809 Patent. Google offers for sale and sells said Google Cloud products with the intent to encourage and facilitate infringing uses of those products in the Northern District of Texas, in the United States, and throughout the world.

88. As a result of Google's infringement of the '809 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

89. In addition, Google's infringing acts and practices have caused and are causing immediate and irreparable harm to Valtrus.

90. Valtrus is informed and believes, and thereon alleges, that Google's infringement of the '809 Patent has been and continues to be willful. As noted above, Google had knowledge of the '809 Patent and its infringement of the '809 Patent no later than April 14, 2021. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Valtrus's patent rights. Thus, Google's infringing actions have been and continue to be consciously wrongful.

FOURTH CLAIM

(Infringement of U.S. Patent No. 7,346,604)

91. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-90 of its Complaint.

92. The '604 Patent, entitled "Method for ranking hypertext search results by analysis of hyperlinks from expert documents and keyword scope," was duly and lawfully issued on March 18, 2008. A true and correct copy of the '604 Patent is attached hereto as Exhibit 4.

93. The '604 Patent names Krishna A. Bharat and George A. Mihaila as co-inventors.

94. The '604 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '604 Patent, including the right to seek damages for any infringement thereof.

95. The '604 Patent "relates generally to software programs and, more specifically, to search engines that search large numbers of hypertext documents." Ex. 4 at 1:9-11.

96. The '604 Patent states that "[a] broad search query can produce a huge result set. This set is hard to rank based on content alone, since the quality and 'authoritativeness' (namely, a measure of how authoritative the page is on the subject) of pages cannot be assessed solely by analyzing their content." *Id.* at 1:31-35.

97. The '604 Patent explains, "[b]y combining relevant out-going links from many experts on the query topic, it is possible to find the pages that are most highly regarded by the community of pages related to the query topic." *Id.* at 3:4-7.

98. The '604 Patent's method for ranking hypertext search results provides an innovative approach to overcoming deficiencies in existing methods of improving the authoritativeness of ranked search results. The '604 Patent describes three such existing methods: "ranking based on human classification," "ranking based on usage information," and "ranking based on connectivity," *id.* at 1:46, 56, 66, and notes that none of these are able to provide authoritative, relevant search results for a large number and variety of queries. *See id.* at 1:44-2:15. The result of the techniques described by the '604 Patent, by contrast, "is to generate


a listing of pages that are highly relevant to the user's query and of high quality," an innovation which "provides a novel and advantageous system and method of searching large numbers of hypertext documents, such as the hypertext documents of the world wide web." *Id.* at 8:27-29, 39-42. The '604 Patent provides specific steps for carrying out this method. For example, in one embodiment, an "expert reverse index" is created by identifying pages which are "about a certain topic and [have] links to many non-affiliated pages on that topic." *Id.* at 4:47-48. Detailed instructions are given for determining whether two hosts are "affiliated" and the degree of a page's expertise, and for "adding keywords in the expert pages to an expert reverse index." *Id.* at 4:57-6:36. This invention improves the performance and efficiency of computers by returning highly relevant and authoritative search results to users while avoiding the high computational overhead and slow response times of other methods. *See id.* at 1:44-2:45.

99. Valtrus is informed and believes, and thereon alleges, that Google has infringed one or more claims of the '604 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed method of analysis of hyperlinks from expert documents in an infringing manner. Google practices every step of at least claim 1 of the '604 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

100. For example, the '604 accused product, Google Search, embodies every limitation of at least claim 1 of the '604 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

101. Google Search practices a computer-implemented method for searching a large number of hypertext documents in accordance with a search query.

102. For example, Google Search “looks through hundreds of billions of webpages” when a user submits a query. *See, e.g.,* How Search works – Organizing information⁷:



103. On information and belief, Google Search practices a method comprising forming a set of expert documents from the set of all hypertext documents crawled without reference to the search query.

104. For example, publicly available information indicates that Google Search uses an algorithm known as “Hilltop” to build a special index of expert documents from the set of all

⁷ <https://www.google.com/search/howsearchworks/how-search-works/organizing-information/>

webpages that Google Search crawlers have indexed. *See, e.g., Aaron Aders, Build Trust and Authority in Google Search, Inc.com*⁸:

To offer an understanding of authority, the Hilltop algorithm uses "a special index of expert documents" as starting points for authority signals. These "expert documents" are simply Web pages that Google has deemed an authoritative Hilltop. When a Hilltop links to a website outside of this special club, that website earns trust and authority. This increased brand authority earns higher search engine rankings.

105. Google Search "systems" form said index of expert documents by determining "which content demonstrates expertise, authoritativeness, and trustworthiness." This determination is made based at least in part on "if other prominent websites link or refer to the content." On information and belief, at least part of this determination occurs without reference to a particular query. *See How Search works – Ranking results – Quality of content*⁹ (red rectangular annotations added):

After identifying relevant content, our systems aim to prioritize those that seem most helpful. To do this, they identify signals that can help determine which content demonstrates expertise, authoritativeness, and trustworthiness.

For example, one of several factors we use to help determine this is understanding if other prominent websites link or refer to the content. This has often proven to be a good sign that the information is well trusted. Aggregated feedback from our Search quality evaluation process is used to further refine how our systems discern the quality of information.

⁸ <https://www.inc.com/aaron-aders/build-trust-and-authority-in-google-search.html>

⁹ <https://www.google.com/search/howsearchworks/how-search-works/ranking-results/#quality>

106. Google Search practices a method comprising ranking the expert documents in accordance with the search query.

¹⁰ (red rectangular annotations added):

How results are automatically generated

With the vast amount of information available, finding what you need would be nearly impossible without some help sorting through it. Google's ranking systems are designed to do just that: sort through hundreds of billions of webpages and other content in our Search index to present the most relevant, useful results in a fraction of a second.

Key factors in your results

To give you the most useful information, Search algorithms look at many factors and signals, including the words of your query, relevance and usability of pages, expertise of sources, and your location and settings. The weight applied to each factor varies

¹⁰ <https://www.google.com/search/howsearchworks/how-search-works/ranking-results/>

After identifying relevant content, our systems aim to prioritize those that seem most helpful. To do this, they identify signals that can help determine which content demonstrates expertise, authoritativeness, and trustworthiness.

109. Google Search practices a method comprising ranking target documents pointed to by the ranked expert documents.

110. For example, Google Search crawlers automatically visit webpages as well as hyperlinks on those pages. Thus, when an expert document is crawled and ranked, Google Search also ranks target documents to which the expert document links. *See, e.g.,* How Search works – Organizing information¹¹ (red rectangular annotation added):

Most of our Search index is built through the work of software known as crawlers. These automatically visit publicly accessible webpages and follow links on those pages, much like you would if you were browsing content on the web. They go from page to page and store information about what they find on these pages and other publicly-accessible content in Google's Search index.

111. Google Search practices a method comprising returning a results list based on the ranked target documents.

112. For example, the results list presented to a Google Search user is based at least in part on whether Google Search has identified a webpage as being a target document linked to by an expert document. Target documents that are linked to by authoritative expert webpages are

¹¹ <https://www.google.com/search/howsearchworks/how-search-works/organizing-information/>

ranked higher than other webpages. *See, e.g., Beyond PageRank: The Hilltop Algorithm*¹² (red rectangular annotations added):

PageRank started becoming slower because the mutually recursive algorithm was inefficient at finding the authority scores the way it was described in class. The Hilltop Algorithm is one of the most important algorithms in search engine applications today. It was introduced by Google in 2003 and continues being the primary foundation of SEO even as other parameters in newer algorithms are added. The primary goal of the Hilltop algorithm was identifying authoritative digital pages, and that information would be used for page ranking. It is essentially a model that worked together with PageRank to identify the pages that are authoritative and expert. Hilltop identifies the authoritative website pages by first selecting the "expert pages." Expert pages mean the pages that are specific to particular topics that linked to a lot of other websites. Once the expert pages are selected, Google goes ahead to identify the highest quality sites to rank.

When it was introduced, Hilltop had a dramatic impact on Search Engine Optimization (SEO). It was established that links that emanated from .edu pages, links from highly relevant websites such as About.com, and links that came from top directories such as Zeal and Yahoo help websites to get better ranking. Note that these high authority pages have continued to increase over time. In fact, getting a link from sites like Yahoo or About.com was known to significantly increase the page ranking.

113. Google has had knowledge of the '604 Patent since at least June 26, 2012. On this date, U.S. Patent No. 8,209,339, which cites the '604 Patent as a reference, was issued with Google Inc. as assignee. The '604 Patent is also cited as a reference by U.S. Patent No. 8,849,817 (issued September 30, 2014), which names Google Inc. as assignee. Furthermore, one of the co-inventors of the '604 Patent, Krishna A. Bharat, is listed as an inventor on two other patents which are assigned to Google: U.S. Patent Nos. 7,359,894 (issued April 15, 2008) and 9,037,575 (issued May 19, 2015). U.S. Patent No. 9,037,575 relates to the ranking of news articles. Bharat, who has been described as a "Google veteran" and "the inventor of Google News," has spent nearly two decades working for Google during the time since the filing of the

¹² <https://blogs.cornell.edu/info2040/2019/10/27/beyond-pagerank-the-hilltop-algorithm/>

'604 Patent application.¹³ Given that two of its own patents cite the '604 Patent as a reference, and that one of the co-inventors of the '604 Patent has since been extensively involved with the development of Google products, Google knew of or was willfully blind to the fact that its conduct was infringing by, at the very latest, June 2012.

114. Valtrus is informed and believes, and thereon alleges, that Google actively, knowingly, and intentionally has induced infringement of the '604 Patent by, for example, offering for public use Google Search with the intent to encourage and facilitate infringing uses of that service in the Northern District of Texas, in the United States, and throughout the world.

115. As a result of Google's infringement of the '604 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

116. In addition, Google's infringing acts and practices have caused immediate and irreparable harm to Valtrus.

117. Valtrus is informed and believes, and thereon alleges, that Google's infringement of the '604 Patent has been willful. As noted above, Google has had knowledge of the '604 Patent and its infringement of the '604 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Valtrus's patent rights. Thus, Google's infringing actions have been consciously wrongful.

¹³ <https://www.cnbc.com/2019/08/22/google-news-inventor-krishna-bharat-returns-after-four-year-hiatus.html>

FIFTH CLAIM

(Infringement of U.S. Patent No. 7,523,454)

118. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-117 of its Complaint.

119. The '454 Patent, entitled "Apparatus and method for routing a transaction to a partitioned server," was duly and lawfully issued on April 21, 2009. A true and correct copy of the '454 Patent is attached hereto as Exhibit 5.

120. The '454 Patent names Francisco J. Romero and Raja Daoud as co-inventors.

121. The '454 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '454 Patent, including the right to seek damages for any infringement thereof.

122. The '454 Patent "pertains to an apparatus and method for routing a transaction to a partition on a server based at least in part on the configuration thereof." Ex. 5 at 1:7-9.

123. The '454 Patent states that "current approaches to load balancing do not recognize the partitions of a server and the various configurations thereof," preventing efficient allocation of server resources between different partitions. *Id.* at 1:56-58.

124. The '454 Patent explains that "the apparatus and method of the invention recognizes and routes [a] transaction to the partition on the server, based at least in part on the configuration of the partition." *Id.* at 2:61-63.

125. The method of the '454 Patent improves the function of a partitioned server by routing transactions to partitions based on the characteristics of the transaction and of the partitions to which they are routed. In contrast to "current approaches to load balancing," the '454 Patent technique does "recognize the partitions of a server and the various configurations thereof," allowing each transaction to be routed to the most suitable partition of a server and

improving the function of the server by enhancing “utilization of the server resources.” *Id.* at 1:56-58, 41.

126. Valtrus is informed and believes, and thereon alleges, that Google has infringed and unless enjoined will continue to infringe one or more claims of the '454 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed method of routing transactions to partitioned servers in an infringing manner. Google practices every step of at least claim 17 of the '454 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

127. For example, the '454 accused product, Google Cloud, embodies every limitation of at least claim 17 of the '454 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

128. Google Cloud practices a method for routing a transaction to a partitioned server.

129. For example, Google Cloud's load balancing feature “redirect[s] traffic away from busy or unavailable” virtual machines (also known as partitions), thereby routing each transaction to a partition on a Google Cloud server. Additionally, these virtual machines serve requests sent to a single address, which forms a unified partitioned server. *See* <https://cloud.google.com/load-balancing/docs>:

Using load balancing for highly available applications

Use load balancing with a regional managed instance group to redirect traffic away from busy or unavailable VM instances, allowing you to provide high availability even during a zonal outage.

130. Google Cloud practices a method comprising identifying a plurality of partitions on said partitioned server based on a network address for each of said plurality of partitions.

131. For example, Google Cloud “[r]esources such as VM instances and load balancers have IP addresses in Google Cloud,” which “enable Google Cloud resources to communicate with other resources” See <https://cloud.google.com/vpc/docs/ip-addresses?hl=en> (red rectangular annotation added):

IP addresses

[Send feedback](#)

Resources such as VM instances and load balancers have IP addresses in Google Cloud. These IP addresses enable Google Cloud resources to communicate with other resources in Google Cloud, in on-premises networks, or on the public internet. This page describes the IP address categorization used by Google Cloud.

132. Google Cloud practices a method comprising determining a configuration of each of said plurality of partitions in response to an event.

133. For example, Google Cloud's load balancing feature determines the "health" of a partition when it distributes new connections among partitions. A partition's "health" is a configuration of the partition. *See, e.g.,* <https://cloud.google.com/load-balancing/docs/network/networklb-backend-service> (yellow highlight added):

a. **The load balancer selects a backend.** The load balancer calculates a hash based on the configured session affinity. It uses this hash to select a backend from among the ones that are *currently healthy* (unless all backends are unhealthy, in which case all backends are considered as long as the failover policy hasn't been configured to drop traffic in this situation). The default session affinity, *NONE*, uses the following hash algorithms:

- **For TCP and unfragmented UDP packets,** a 5-tuple hash of the packet's source IP address, source port, destination IP address, destination port, and the protocol

134. Google Cloud practices a method comprising determining at least one characteristic of said transaction.

135. For example, Google Cloud's load balancing feature determines the packet's source IP address and distributes new connections to partitions based on that packet's source IP address, which is a characteristic of an HTTPS transaction. *See, e.g., id.* (yellow highlight added):

a. **The load balancer selects a backend.** The load balancer calculates a hash based on the configured session affinity. It uses this hash to select a backend from among the ones that are *currently healthy* (unless all backends are unhealthy, in which case all backends are considered as long as the failover policy hasn't been configured to drop traffic in this situation). The default session affinity, *NONE*, uses the following hash algorithms:

- **For TCP and unfragmented UDP packets,** a 5-tuple hash of the *packet's source IP address*, source port, destination IP address, destination port, and the protocol

136. Google Cloud practices a method comprising routing said transaction to one of said plurality of partitions based on said configurations of said plurality of partitions and based on said at least one characteristic of said transaction.

137. After determining the configuration of each of a plurality of partitions in response to an event and the characteristic(s) of a transaction (such as its IP address), as described above, Google Cloud's load balancing feature routes the transaction to one of said plurality of partitions.

138. Google had knowledge of the '454 Patent no later than May 27, 2020. Raja Daoud, one of the co-inventors of the '454 Patent, is listed as a co-inventor on European Patent Office Patent No. 3,356,961 (issued May 27, 2020) and U.S. Patent Application Publication No. 2021/0184862 (published June 17, 2021), both of which name Google as assignee. Furthermore, on April 14, 2021, Valtrus sent to Google a notice letter with an exemplary list of Valtrus patents, which included the '454 Patent, that were infringed by specific exemplary Google products. Google therefore knew of or was willfully blind to its infringement no later than April 14, 2021.

139. Valtrus is informed and believes, and thereon alleges, that Google actively, knowingly, and intentionally has induced infringement by, for example, using, selling, and offering for sale Google Cloud products which use the process claimed by the '454 Patent. Google offers for sale and sells said Google Cloud products with the intent to encourage and facilitate infringing uses of those products in the Northern District of Texas, in the United States, and throughout the world.

140. As a result of Google's infringement of the '454 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

141. In addition, Google's infringing acts and practices have caused and are causing immediate and irreparable harm to Valtrus.

142. Valtrus is informed and believes, and thereon alleges, that Google's infringement of the '454 Patent has been and continues to be willful. As noted above, Google has had knowledge of the '454 Patent and its infringement of the '454 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Valtrus's patent rights. Thus, Google's infringing actions have been and continue to be consciously wrongful.

SIXTH CLAIM

(Infringement of U.S. Patent No. 7,748,005)

143. Valtrus re-alleges and incorporates herein by reference Paragraphs 1-142 of its Complaint.

144. The '005 Patent, entitled "System and method for allocating a plurality of resources between a plurality of computing domains," was duly and lawfully issued on June 29, 2010. A true and correct copy of the '005 Patent is attached hereto as Exhibit 6.

145. The '005 Patent names Francisco Romero, Cliff McCarthy, and Scott Rhine as co-inventors.

146. The '005 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '005 Patent, including the right to seek damages for any infringement thereof.

147. The '005 Patent "is generally related to allocating a plurality of resources between a plurality of computing domains." Ex. 6 at 1:19-21.

148. The '005 Patent states that "[c]omputer systems inherently have limited resources, particularly CPU resources," which "must be allocated among the different applications operating within the system. A known allocation mechanism . . . is a system known as a Process

Resource Manager (PRM). It is used to partition the CPU resource and various other resources among the different applications.” *Id.* at 1:25-32.

149. The ’005 Patent explains that “virtualization software creates virtual resources as software constructs. The virtual resources are then assigned to virtual machines. Specifically, the virtual resources are used to execute ‘guest’ operating systems that execute on top of the host operating system. The guest operating systems are then used to execute applications. The assignment of the virtual resources to the virtual machines thereby allocates resources between the respective applications.” *Id.* at 1:41-48.

150. The method of the ’005 Patent improves the function of partitioned virtual machines by dynamically reallocating the inherently limited resources of a computer system “in response to received requests for additional resources according to service level parameters.” *Id.* at 2:33-34. This helps to ensure that said limited resources are efficiently allocated and reallocated to computing domains where they are needed most, which improves the computer’s performance by ensuring its access to adequate CPU and other resources.

151. Valtrus is informed and believes, and thereon alleges, that Google has infringed one or more claims of the ’005 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, methods of allocating a plurality of resources between computing domains. Google practices every step of at least claim 8 of the ’005 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.


152. For example, the ’005 accused product, Cloud Run, embodies every limitation of at least claim 8 of the ’005 Patent, literally or under the doctrine of equivalents, as set forth

below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

153. Google's Cloud Run product practices a method comprising the following elements.

154. Google's Cloud Run product practices a method comprising creating a plurality of computing domains.

155. For example, Google Cloud Platforms offers Cloud Run to create and manage a plurality of computing domains for services that run one or more "container" instances. *See, e.g.,* <https://cloud.google.com/run/docs> (blue highlight added):



156. Google's Cloud Run product practices a method comprising allocating a plurality of resources between said plurality of computing domains by a first manager process.

157. For example, Cloud Run allocates a plurality of resources, such as CPU and memory, between a plurality of computing domains (or "containers") running applications on Google Cloud infrastructure. These resources are allocated according to a default manager process. *See* <https://cloud.google.com/run/docs/reference/container-contract>:

Container instance resources

CPU

Cloud Run allocates 1 vCPU per container instance by default, but this can be changed. To specify a different vCPU value, refer to the documentation on [CPU limits](#).

A vCPU is implemented as an abstraction of underlying hardware to provide the approximate equivalent CPU time of a single hardware hyper-thread on variable CPU platforms. The container instance may be executed on multiple cores simultaneously.

You can specify CPU to be always allocated during the life of the instance or to be only allocated during container instance startup and request processing. Refer to [CPU allocation](#) for details.

If you have configured a number of [minimum instances](#), these instances are also subject to the CPU allocation configuration.


Memory

Each Cloud Run container instance by default gets 512 MiB of memory. You can change this by [configuring memory limits](#), up to a maximum of 16 GiB.


158. Google's Cloud Run product practices a method comprising executing at least one application, a second manager process, and a performance monitor process in each of said plurality of computing domains.

159. For example, Cloud Run customers are able to write code for a web application and "push" this code to Google Cloud to create a container instance. This code is then executed on Google Cloud resources to create a "fully managed web application." *See*

<https://cloud.google.com/blog/topics/developers-practitioners/cloud-run-story-serverless-containers> (blue highlight added):



160. When said container application has been deployed, a second manager process allows the initial deployment to be updated and modified automatically or by the user. For example, a user can “specify an instance to kept idle or ‘warm’” if so desired. *See* <https://cloud.google.com/run/docs/about-instance-autoscaling> (blue highlight added):



161. Each Cloud Run computing domain is “automatically integrated with Cloud Monitoring,” which provides “performance monitoring, metrics, and uptime checks.” *See* <https://cloud.google.com/run/docs/monitoring> (blue highlight added):

Monitoring Health and Performance

[Send feedback](#)

[Cloud Monitoring](#)  provides Cloud Run performance monitoring, [metrics](#) , and [uptime checks](#) , along with [alerts](#)  to send notifications when certain metric thresholds are exceeded. [Google Cloud's operations suite pricing](#)  applies, which means there is no charge for metrics on the fully managed version of Cloud Run. Note that you can also use [Cloud Monitoring custom metrics](#).

Cloud Run is automatically integrated with Cloud Monitoring with no setup or configuration required. This means that metrics of your Cloud Run services are captured automatically when they are running.

162. Google's Cloud Run product practices a method wherein said second manager process maintains a list comprising a plurality of application priority levels for said at least one application and an indication of a quantity of said plurality of resources needed to meet said at least one level of said application priority levels.

163. For example, the Cloud Run second manager process maintains a list of priority levels for a given application, including, *inter alia*, request timeout settings. *See, e.g.*, <https://cloud.google.com/run/docs/configuring/request-timeout>:

Setting request timeout

[Send feedback](#)

The request timeout setting specifies the time within which a response must be returned by services deployed to Cloud Run. If a response isn't returned within the time specified, the request ends and error 504 is returned.

The timeout is set by default to 5 minutes and can be extended up to 60 minutes.

164. The Cloud Run second manager process also maintains an indication of a quantity of resources needed to meet at least one level of said application priority levels. For example, Cloud Run container instances are allocated specific quantities of CPU and memory resources by default. These quantities can be adjusted based on the user's needs. *See* <https://cloud.google.com/run/docs/configuring/cpu-allocation> (blue highlight added):

By default, Cloud Run container instances are only allocated CPU during request processing and container startup. You can change this behavior so CPU is always allocated and available even when there are no incoming requests. Setting the CPU to be always allocated can be useful for running background tasks and other asynchronous processing tasks.

See also <https://cloud.google.com/run/docs/configuring/memory-limits> (blue highlight added):

Setting and updating memory limits

Any configuration change leads to the creation of a new revision. Subsequent revisions will also automatically get this configuration setting unless you make explicit updates to change it.

By default, the memory allocated to each container instance of a revision is 512 MiB.

A minimum of 2 vCPUs is required to set a memory limit higher than 4GiB.

You can set memory limits using the Cloud Console, the gcloud command line, or a YAML file when you [create a new service](#) or [deploy a new revision](#):

165. Google's Cloud Run product practices a method wherein said performance monitor process generates performance data related to said at least one application and said second manager process requests additional resources from the first manager process in response to analysis of said performance data in view of at least one service level parameter.

166. For example, the Cloud Monitoring API defines certain metrics for the Cloud Run service and generates data related to the application(s) running in the container(s), such as CPU allocation and utilization. *See, e.g.,* https://cloud.google.com/monitoring/api/metrics_gcp#gcp-run:

Metric type Launch stage	Display name
Kind, Type, Unit Monitored resources	Description Labels
<code>container/billable_instance_time</code> GA Billable Instance Time	
DELTA, DOUBLE, s <code>cloud_run_revision</code>	Billable time aggregated from all container instances of the revision. For a given container instance, billable time occurs when the container instance is starting or at least one request is being processed. Billable time is rounded up to the nearest 100 milliseconds. Examples: If a revision with 2 container instances has been continuously serving traffic in the last minute, the value is 2s/s with the default "rate" aligner. If a single request lasting 30ms was received by a revision in the past minute, it is rounded up to 100ms and averaged to 1.7ms/s over the minute with the default "rate" aligner. Sampled every 60 seconds. After sampling, data is not visible for up to 180 seconds.
<code>container/cpu/allocation_time</code> GA Container CPU Allocation	
DELTA, DOUBLE, s{CPU} <code>cloud_run_revision</code>	Container CPU allocation of the revision in seconds. Sampled every 60 seconds. After sampling, data is not visible for up to 180 seconds.
<code>container/cpu/utilizations</code> GA Container CPU Utilization	
DELTA, DISTRIBUTION, 10^2. % <code>cloud_run_revision</code>	Container CPU utilization distribution across all container instances of the revision. Sampled every 60 seconds. After sampling, data is not visible for up to 60 seconds.
<code>container/instance_count</code> GA Instance Count	
GAUGE, INT64, 1 <code>cloud_run_revision</code>	Number of container instances that exist, broken down by state. Sampled every 60 seconds. After sampling, data is not visible for up to 180 seconds. state: Whether a container is "active" or "idle".

167. Furthermore, Cloud Run’s second manager process requests additional resources (i.e., “scales”) in response to the analysis of the performance monitor described in paragraph 165 *supra* (i.e., “rate of incoming requests”) and in view of at least one service level parameter such as maximum/minimum instances and/or concurrency. *See* <https://cloud.google.com/run/docs/about-instance-autoscaling> (red rectangular annotations added):

About container instance autoscaling

[Send feedback](#)

In Cloud Run, each [revision](#) is [automatically scaled](#) to the number of container instances needed to handle all incoming requests. When a revision does not receive any traffic, by default it is scaled in to zero container instances. However, if desired, you can change this default to specify an instance to be kept idle or "warm" using the [minimum instances](#) setting.

In addition to the [rate of incoming requests](#), the number of instances scheduled is impacted by:

- The CPU utilization of existing instances (Targeting to keep scheduled instances to a 60% CPU utilization)
- The [maximum concurrency setting](#)
- The [maximum number of container instances setting](#)
- The [minimum number of container instances setting](#)

In some cases you may want to limit the total number of container instances that can be started, for cost control reasons, or for better compatibility with other resources used by your service. For example, your Cloud Run service might interact with a database that can only handle a certain number of concurrent open connections.

168. Google's Cloud Run product practices a method comprising dynamically reallocating said plurality of resources between said plurality of computing domains by the first manager process in response to received requests for additional resources according to service level parameters.

169. For example, "in response to incoming requests," Cloud Run dynamically reallocates resources (i.e., "automatically scale[s]") to "a certain number of container instances." See <https://cloud.google.com/run/docs/reference/container-contract> (red annotation added):

Container instance lifecycle

[In response to incoming requests, a service is automatically scaled to a certain number of container instances](#), each of which runs the deployed container image.

When a revision does not receive any traffic, it is scaled in to the [minimum number of container instances configured](#) (zero by default).

170. As a result of Google's infringement of the '005 Patent, Valtrus has been damaged. Valtrus is entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

PRAYER FOR RELIEF

WHEREFORE, Valtrus prays for judgment against Google as follows:

- A. That Google has infringed each of the Asserted Patents, and unless enjoined will continue to infringe one or more of the Asserted Patents;
- B. That Google has willfully infringed one or more of the Asserted Patents;
- C. That Google pay Valtrus damages adequate to compensate Valtrus for Google's past infringement of each of the Asserted Patents, and present and future infringement of applicable Asserted Patents, together with interest and costs under 35 U.S.C. § 284;
- D. That Google be ordered to pay prejudgment and post-judgment interest on the damages assessed;
- E. That Google pay Valtrus enhanced damages pursuant to 35 U.S.C. § 284;
- F. That Google be ordered to pay supplemental damages to Valtrus, including interest, with an accounting, as needed;
- G. That Google be enjoined from infringing the applicable Asserted Patents, or if its infringement is not enjoined, that Google be ordered to pay ongoing royalties to Valtrus for any post-judgment infringement of the Asserted Patents;
- H. That this is an exceptional case under 35 U.S.C. § 285, and that Google pay Valtrus's attorneys' fees and costs in this action; and
- I. That Valtrus be awarded such other and further relief, including equitable relief, as this Court deems just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), Valtrus hereby demands a trial by jury on all issues triable to a jury.

January 10, 2022

By: /s/ Jon B Hyland
Jon B. Hyland (State Bar No. 24046131)
HILGERS GRABEN
10000 N. Central Expressway, Suite 400
Dallas, Texas 75231
Phone: (972) 645-7728
E-mail: jhyland@hilgersgraben.com

COUNSEL FOR PLAINTIFF