

1 Douglas G. Muehlhauser (State Bar No. 179,495)  
doug.muehlhauser@knobbe.com  
2 Alan G. Laquer (State Bar No. 259,257)  
alan.laquer@knobbe.com  
3 KNOBBE, MARTENS, OLSON & BEAR, LLP  
4 2040 Main Street, Fourteenth Floor  
Irvine, CA 92614  
5 Telephone: 949-760-0404  
Facsimile: 949-760-9502  
6

7 Attorneys for Plaintiff Pepperdata, Inc.  
8

9 IN THE UNITED STATES DISTRICT COURT  
10 FOR THE NORTHERN DISTRICT OF CALIFORNIA

11  
12 PEPPERDATA, INC.,  
13 Plaintiff,  
14 v.  
15 YAHOO! INC.,  
16 Defendant.

Civil Action No. 5:16-CV-7211-EJD

**FIRST AMENDED COMPLAINT  
FOR INFRINGEMENT OF  
U.S. PATENT NOS. 8,849,891  
AND 9,325,593**

**and**

**DEMAND FOR JURY TRIAL**

1 Plaintiff Pepperdata, Inc. (“Pepperdata”) brings this complaint for patent infringement  
2 against Yahoo! Inc. (“Yahoo”) and alleges as follows:

3 **I. NATURE OF THE ACTION**

4 1. This is an action for patent infringement arising under the patent laws of the  
5 United States, 35 U.S.C. § 100 *et seq.*

6 **II. PARTIES**

7 2. Pepperdata is a Delaware corporation with its principal place of business at  
8 19409 Stevens Creek Boulevard, Suite 260, Cupertino, California 95014.

9 3. On information and belief, Yahoo is a Delaware corporation with its principal  
10 place of business at 701 First Avenue, Sunnyvale, California 94089.

11 **III. JURISDICTION**

12 4. This Court has subject-matter jurisdiction over this patent infringement action  
13 under 28 U.S.C. §§ 1331 and 1338(a).

14 5. On information and belief, Yahoo has its principal place of business in this  
15 judicial district and is subject to general personal jurisdiction here.

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

18 **IV. FACTS**

19 7. Pepperdata was founded in 2012 and provides its customers with its innovative  
20 product that improves the performance of Hadoop distributed computing clusters, for  
21 example by adding functionality to a Hadoop cluster to dynamically overcommit resources in  
22 that cluster.

23 8. Pepperdata has obtained patents to protect its business investments in  
24 researching and developing its innovative solutions for improving the performance of Hadoop  
25 environments, including dynamically overcommitting resources in a Hadoop cluster.

26 9. On September 30, 2014, the United States Patent and Trademark Office  
27 (“PTO”) issued U.S. Patent No. 8,849,891 (“the ’891 patent”), titled “Systems, Methods, and  
28 Devices for Dynamic Resource Monitoring and Allocation in a Cluster System.” A copy of

1 the '891 patent is provide as Exhibit A to this Complaint.

2 10. On April 26, 2016, the PTO issued U.S. Patent No. 9,325,593 (“the '593  
3 patent”) titled “Systems, Methods, and Devices for Dynamic Resource Monitoring and  
4 Allocation in a Cluster System.” A copy of the '593 patent is provided as Exhibit B to this  
5 Complaint.

6 11. Pepperdata owns the '891 and '593 patents.

7 12. Pepperdata has marked the products that it has manufactured and sold under  
8 the '891 and '593 patents in a manner that complies with 35 U.S.C. § 287(a).

9 13. Yahoo uses Hadoop clusters within the United States.

10 14. Yahoo made the following statements in a November 18, 2016 Hadoop group  
11 tumblr post at [http://yahoohadoop.tumblr.com/post/153336735536/10-years-of-hadoop-and-](http://yahoohadoop.tumblr.com/post/153336735536/10-years-of-hadoop-and-its-israeli-pioneering)  
12 [its-israeli-pioneering](http://yahoohadoop.tumblr.com/post/153336735536/10-years-of-hadoop-and-its-israeli-pioneering)

- 13 • “The Apache Hadoop technology suite is the engine behind the Big  
14 Data revolution that has been transforming multiple industries over  
15 the last decade.”
- 16 • “These days, Yahoo is the largest Hadoop deployment in the  
17 industry. We run tens of thousands of Hadoop machines in our  
18 datacenters and manage more than 600 petabytes of data. Our  
19 products use Hadoop in a variety of ways that reflect a wealth of data  
20 processing patterns.”
- 21 • “Yahoo’s commitment to Hadoop goes far beyond operating the  
22 technology at Web scale. The company’s engineers and scientists  
23 make contributions to both entrenched and incubating Hadoop  
24 projects.”
- 25 • “Our team launched approximately three years ago. Collectively we  
26 add many years of experience to Yahoo and the Hadoop community  
27 in distributed computing research and development.
- 28 • “Our researchers regularly present their innovations at leading

1 industrial conferences (Hadoop Summit and HBaseCon), as well as  
2 at top academic venues.”

3 15. Yahoo made a software patch identified as YARN-5202, titled “Dynamic  
4 Overcommit of Node Resources.”

5 16. Nathan Roberts, a Yahoo software architect, is the assignee of the JIRA for  
6 YARN-5202, accessible at <https://issues.apache.org/jira/browse/YARN-5202>.

7 17. Yahoo made the YARN-5202 software patch available for download on the  
8 YARN-5202 JIRA. For example, YARN-5202 was available for download from its JIRA on  
9 August 4, 2016.

10 18. Upon information and belief, Yahoo made the YARN-5202 software patch in  
11 the United States in 2016.

12 19. Yahoo’s YARN-5202 software patch adds to a Hadoop cluster functionality  
13 for dynamic overcommitting node resources.

14 20. The YARN-5202 JIRA explains that “[t]his Jira is to present a proof-of-  
15 concept implementation (collaboration between Jason Lowe and myself) of a dynamic over-  
16 commit implementation in YARN.”

17 21. Jason Lowe has the title “Senior Principal Engineer, Hadoop” at Yahoo.

18 22. On June 28, 2016, Mr. Lowe gave a presentation, on behalf of Yahoo, titled  
19 “Investigating the Effects of Overcommitting YARN Resources” at the 2016 Hadoop Summit  
20 in San Jose, California (“Yahoo’s YARN-5202 Presentation”).

21 23. Over 3,000 people attended the 2016 Hadoop Summit in San Jose, California.

22 24. As of December 15, 2016, a video of Yahoo’s YARN-5202 Presentation is  
23 available at <https://www.youtube.com/watch?v=hILD2g9putc>, and this video of Yahoo’s  
24 YARN-5202 Presentation was published and made available for public viewing at least by  
25 June 29, 2016.

26 25. A transcript of Yahoo’s YARN-5202 Presentation is provided as Exhibit C to  
27 this Complaint.

28 26. The YARN-5202 JIRA shows that by June 14, 2016, Yahoo had uploaded the

1 YARN-5202 patch to make it available to the public for free download, installation and use.  
2 A copy of the YARN-5202 JIRA, printed on July 14, 2016, is attached hereto as Exhibit D.

3 27. The YARN-5202 JIRA describes the YARN-5202 patch as an “improvement”  
4 having a “major” priority.

5 28. Upon information and belief, Yahoo has used the YARN-5202 software patch  
6 on Hadoop clusters that it uses in the United States.

7 29. Mr. Lowe explained in Yahoo’s YARN-5202 Presentation that Yahoo has  
8 used the YARN-5202 software patch on Yahoo’s main research cluster, which runs Hadoop  
9 YARN.

10 30. Mr. Lowe explained in Yahoo’s YARN-5202 Presentation that Yahoo has  
11 enjoyed a “significant improvement” in the performance of its Hadoop cluster as a result of  
12 the dynamic overcommit functionality in the YARN-5202 software patch. For example, Mr.  
13 Lowe explained that:

14 before [adding dynamic overcommit] utilization was hovering  
15 like around 40/50% now it’s hovering – especially during the  
16 peg thing – it’s hovering closer to 70/80% so that’s a  
17 significant improvement for us – it’s like close to 40/50% CPU  
18 utilization improvement. And so that’s been a big deal for us.  
19 This like down at the bottom you can see that in terms of the  
20 YARN size of the cluster reported it’s like 50% bigger which is  
21 like taking your cluster and putting half again as many nodes  
22 on the cluster without actually buying those nodes. So it’s a  
23 really big deal for us.

24 31. On August 4, 2016, Pepperdata provided to Yahoo written notice that Yahoo  
25 infringes the ’891 and ’593 patents. A copy of Pepperdata’s August 4, 2016 letter to Yahoo  
26 is attached hereto as Exhibit E.

27 32. Pepperdata’s August 4, 2016 letter to Yahoo included claim charts showing,  
28 by detailed and exemplary explanation, how Yahoo’s use of its YARN-5202 patch in a

1 Hadoop computer cluster falls within the scope of representative claims of the '891 and '593  
 2 patents.

3 33. For example, the claim chart provided to Yahoo on August 4, 2016, includes  
 4 the following detailed exemplary explanation of how Yahoo's use of its YARN-5202 patch in  
 5 a Hadoop computer cluster falls within the scope of claim 1 of the '891 patent:

6 U.S. Patent No. 8,849,891 7 Claim 1	Yahoo's Hadoop Clusters with Yahoo's Dynamic Overcommit of Node Resources
8 1. A computer cluster comprising: 9 a management computing device 10 comprising a supervisor controller 11 configured to coordinate processing of 12 a plurality of sub-jobs for a plurality of 13 overall jobs; 14 15 16 17 18 19 20 21	Each of Yahoo's Hadoop computer clusters comprises a management computing device running a ResourceManager, which comprises a supervisor controller configured to coordinate processing of a plurality of sub-jobs for a plurality of overall jobs. For example, the Hadoop MapReduce Tutorial1 explains "The MapReduce framework consists of a single master ResourceManager, one slave NodeManager per cluster-node, and MRAppMaster per application." As another example, Yahoo's jira2 for its modified Hadoop software identifies its components as "nodemanager, resourcemanager."
22 a plurality of computer system nodes 23 configured to communicate with the 24 management computing device, and to 25 perform processing of received 26 sub-jobs, the computing system nodes 27 each comprising: 28	Each of Yahoo's Hadoop computer clusters comprises a plurality of computer system nodes configured to communicate with the management computing device, and to perform processing of received sub-jobs.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

<p>one or more processors configured to perform computing processes on received sub-jobs;</p>	<p>Each of the computing system nodes comprises one or more processors configured to perform computing processes on received sub-jobs.</p> <p>For example, the Hadoop NodeManager Overview explains<sup>3</sup> that “The NodeManager is responsible for launching and managing containers on a node. Containers execute tasks as specified by the AppMaster.”</p>
<p>an agent controller comprising:</p>	<p>Each of the computing systems nodes comprises a NodeManager which comprises an agent controller.</p>
<p>a monitoring interface configured to monitor utilization by sub-jobs of system resources of a first computing system node; and</p>	<p>The agent controller of the NodeManager comprises a monitoring interface configured to monitor utilization by sub-jobs of system resources of a first computing system node.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture<sup>4</sup> describes: “The NodeManager is the per-machine framework agent who is responsible for containers, monitoring their resource usage (cpu, memory, disk, network) and reporting the same to the ResourceManager/Scheduler.”</p>
<p>a reporting controller configured to transmit the monitored system resources utilization to the supervisor controller in substantially real-time;</p>	<p>The agent controller of the NodeManager comprises a reporting controller configured to transmit the monitored system resources utilization to the supervisor controller in substantially real-time.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture describes: “The NodeManager is the per-machine framework agent who is responsible for containers,</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

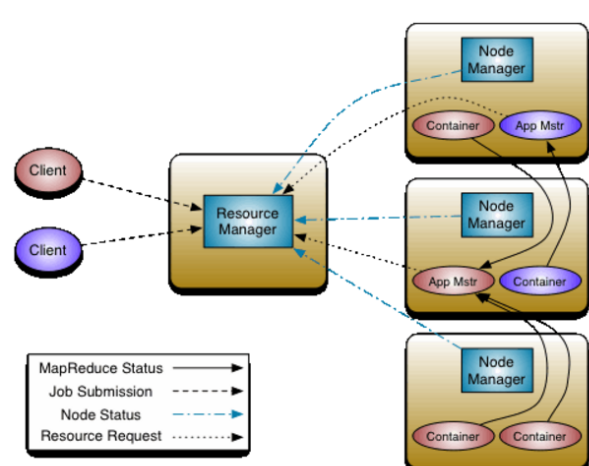
	<p>monitoring their resource usage (cpu, memory, disk, network) and reporting the same to the ResourceManager/Scheduler.”</p> <p>Yahoo explained in its dynamic overcommit presentation<sup>5</sup> that its NodeManager reports this information in its heartbeat. (See at 10 minute mark.)</p>
<p>wherein the supervisor controller is configured to assign an additional sub-job to the first computing system node based on determining that the utilization of at least one system resource of the first computing system node is below a threshold level, the determining based on the monitored system resources utilization transmitted from the reporting controller to the supervisor controller;</p>	<p>The supervisor controller of the ResourceManager is configured to assign an additional sub-job to the first computing system node based on determining that the utilization of at least one system resource of the first computing system node is below a threshold level, the determining based on the monitored system resource utilization transmitted from the reporting controller of the NodeManager to the supervisor controller of the ResourceManager.</p> <p>For example, the ResourceManager determines that the memory utilization reported by NodeManager is below the corresponding low water mark, as reflected in the “conf” variable, and therefore assigns one or more additional sub-jobs to the first computing system node by incrementing the containers allocated on that node, such as through the memIncrement = getConfInt(conf,...) function call at line 1249 of Yahoo’s dynamic overcommit patch.</p> <p>As another example, the ResourceManager</p>



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

	<p>determines that the virtual core utilization reported by NodeManager is below the corresponding low water mark, as reflected in the “conf” variable, and therefore assigns one or more additional sub-jobs to the first computing system node by incrementing the containers allocated on that node, such as through the <code>vcoreIncrement = getConfInt(conf,...)</code> function call at line 1272 of Yahoo’s dynamic overcommit patch.</p> <p>Yahoo explained in its dynamic overcommit presentation: “How does the ResourceManager do that scaling? We have a high-level watermark and a low watermark... As long as the node keeps reporting lower than the low watermark utilization we’ll increment over time the size of that node inside the ResourceManager so that it’ll keep allocating more resources inside that node.” (See at 12 minute mark.)</p> <p>As another example, in Yahoo’s dynamic overcommit patch, the <code>updateTotalResource</code> function is described as “Adjust overcommit metrics... The amount to overcommit will be re-calculated on next node heartbeat.”</p>
<p>wherein the at least one system resource of the first computing system node is a first electronic random access memory capacity,</p>	<p>As described above in the context of the <code>memIncrement = getConfInt(conf,...)</code> function call, at least one system resource of the first computing system node is a first electronic random access</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

	<p>memory capacity.</p> <p>As another example, in Yahoo’s dynamic overcommit patch, Yahoo’s source code comments explain that RM_OVERCOMMIT_MEM_LWM corresponds to:</p> <p>“Low water memory utilization mark for overcommit. If the node's memory utilization is below this value then the scheduler will try to maximize the memory overcommit.”</p>
<p>wherein the supervisor controller is configured to monitor a second electronic random access memory capacity of a second computing system node,</p>	<p>The supervisor controller of the ResourceManager is configured to monitor a second electronic random access memory capacity of a second computing system node. For example, the ResourceManager receives memory utilization information from the node’s NodeManager.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture<sup>6</sup> illustrates that the architecture includes multiple nodes:</p>  <p>The diagram illustrates the Hadoop 2.7.2 YARN Architecture. It shows a central Resource Manager (RM) box connected to three Node Manager (NM) boxes. Each NM box contains an App Master (App Mstr) and one or more Containers. Two Client boxes are shown on the left, connected to the RM. A legend at the bottom left defines the types of arrows: solid for MapReduce Status, dashed for Job Submission, dashed for Node Status, and dotted for Resource Request. Arrows indicate the flow of information: Clients submit jobs to the RM; the RM sends job submissions to Node Managers; Node Managers report node status to the RM; and Node Managers request resources from the RM.</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

<p>wherein the assigning by the supervisor controller of the additional sub-job comprises assigning the additional sub-job to the first computing system node based on determining that utilization of the first electronic random access memory capacity is below the threshold level,</p>	<p>The supervisor controller of the ResourceManager assigns the additional sub-job to the first computing system node based on determining that utilization of the first electronic random access memory capacity is below the corresponding low watermark threshold level.</p> <p>For example, the ResourceManager determines that the memory utilization reported by NodeManager is below the corresponding low water mark, as reflected in the “conf” variable, and therefore assigns one or more additional sub-jobs to the first computing system node by incrementing the containers allocated on that node, such as through the memIncrement = getConfInt(conf,...) function call at line 1249 of Yahoo’s dynamic overcommit patch.</p> <p>Yahoo explained in its dynamic overcommit presentation: “How does the ResourceManager do that scaling? We have a high-level watermark and a low watermark... As long as the node keeps reporting lower than the low watermark utilization we’ll increment over time the size of that node inside the ResourceManager so that it’ll keep allocating more resources inside that node.” (See at 12 minute mark.)</p>
<p>wherein the supervisor controller is configured to prevent assignment of</p>	<p>The supervisor controller of the ResourceManager is configured to prevent assignment of additional</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

<p>additional sub-jobs to the second computing system node based on determining that utilization of the second electronic random access memory capacity is at or above a threshold value,</p>	<p>sub-jobs to the second computing system node based on determining that utilization of the second electronic random access memory capacity is at or above a high watermark threshold value.</p> <p>For example, at line 453 of Yahoo’s dynamic overcommit patch, Yahoo’s code checks whether the value of the variable “pmemUsedPercent” is greater than or equal to the highWaterMark and, if it is, sets needToTrim = true. Because needToTrim is set to true, the function does not return at line 486 and instead goes on to perform trim functions.</p>
<p>wherein the additional sub-job requires utilization of the first electronic random access memory capacity that is unused on the first computing system node.</p>	<p>The additional sub-jobs assigned to the require utilization of the first electronic random access memory capacity that is unused on the first computing resource code.</p> <p>For example, Yahoo explained in its dynamic overcommit presentation that the ResourceManager scales overcommit of resources so that the node utilizes the resources (particularly RAM) to perform the sub-jobs corresponding to its allocated containers. (See at 12 minute mark.)</p>

1 34. Also, for example, the claim chart provided to Yahoo on August 4, 2016,  
 2 includes the following detailed exemplary explanation of how Yahoo’s use of its YARN-  
 3 5202 patch in a Hadoop computer cluster falls within the scope of claim 1 of the ’593 patent:

4 U.S. Patent No. 9,325,593 5 Claim 1	Yahoo’s Hadoop Clusters with Yahoo’s Dynamic Overcommit of Node Resources
6 1. A hadoop computer cluster 7 comprising:	
8 one or more processors of a master 9 node, wherein the master node 10 comprises a supervisor controller; 11 12 13 14 15 16 17	Each of Yahoo’s Hadoop computer clusters comprises one or more processors of a master node, wherein the master node comprises a ResourceManager, which comprises a supervisor controller. For example, the Hadoop MapReduce Tutorial1 explains “The MapReduce framework consists of a single master ResourceManager, one slave NodeManager per cluster-node, and MRAppMaster per application.”
18 one or more processors of a plurality 19 of computing system nodes, the one or 20 more processors of the plurality of 21 computing system nodes configured to 22 perform computing processes on 23 received sub-jobs, wherein each 24 computing system node comprises an 25 agent controller; 26 27 28	Each of Yahoo’s Hadoop computer clusters comprises one or more processors of a plurality of computing system nodes, the one or more processors of the plurality of computing system nodes configured to perform computing processes on received sub-jobs, wherein each computing system node comprises a NodeManager which comprises an agent controller. For example, the Hadoop MapReduce Tutorial explains “The MapReduce framework consists of a

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

	<p>single master ResourceManager, one slave NodeManager per cluster-node, and MRAppMaster per application.”</p> <p>For example, the Hadoop NodeManager Overview explains<sup>2</sup> that “The NodeManager is responsible for launching and managing containers on a node. Containers execute tasks as specified by the AppMaster.”</p>
<p>wherein each agent controller comprises:</p>	<p>Each agent controller comprises...</p>
<p>a monitoring interface configured to monitor system resources utilization by sub-jobs of its respective computing system node; and</p>	<p>Each agent controller of each NodeManger comprises a monitoring interface configured to monitor system resources utilization by sub-jobs of its respective computing system node.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture<sup>3</sup> describes: “The NodeManager is the per-machine framework agent who is responsible for containers, monitoring their resource usage (cpu, memory, disk, network) and reporting the same to the ResourceManager/Scheduler.”</p>
<p>a reporting controller configured to transmit the monitored system resources utilization to the supervisor controller in substantially realtime;</p>	<p>Each agent controller of each NodeManager comprises a reporting controller configured to transmit the monitored system resource utilization to the supervisor controller of the ResourceManager in substantially real-time.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture describes: “The NodeManager is the per-machine</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

	<p>framework agent who is responsible for containers, monitoring their resource usage (cpu, memory, disk, network) and reporting the same to the ResourceManager/Scheduler.”</p> <p>Yahoo explained in its dynamic overcommit presentation<sup>4</sup> that its NodeManager reports this information in its heartbeat. (See at 10 minute mark.)</p>
<p>wherein the supervisor controller is configured to assign an additional sub-job to a first computing system node based on determining that the utilization of a first electronic random access memory capacity of the first computing system node is below a threshold level, the determining based on the monitored system resources utilization transmitted from the reporting controller of the first computing system node to the supervisor controller,</p>	<p>The supervisor controller of the ResourceManager is configured to assign an additional sub-job to a first computing system node based on determining that the utilization of a first electronic random access memory capacity of the first computing system node is below a low watermark threshold level, the determination based on the monitored system resources utilization transmitted from the reporting controller of the NodeManager of the first computing system node to the supervisor controller of the ResourceManager.</p> <p>For example, the ResourceManager determines that the memory utilization reported by NodeManager is below the corresponding low water mark, as reflected in the “conf” variable, and therefore assigns one or more additional sub-jobs to the first computing system node by incrementing the containers allocated on that node, such as through the memIncrement = getConfInt(conf,...) function</p>

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

call at line 1249 of Yahoo’s dynamic overcommit patch.

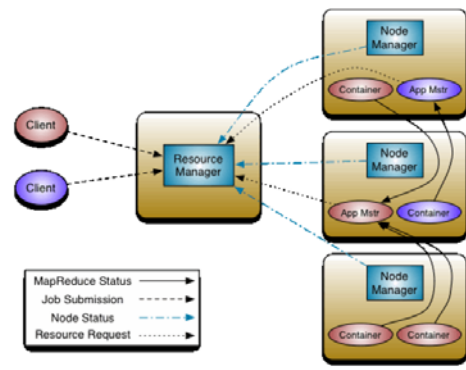
Yahoo explained in its dynamic overcommit presentation: “How does the ResourceManager do that scaling? We have a high-level watermark and a low watermark... As long as the node keeps reporting lower than the low watermark utilization we’ll increment over time the size of that node inside the ResourceManager so that it’ll keep allocating more resources inside that node.” (See at 12 minute mark.)

As another example, in Yahoo’s dynamic overcommit patch, the updateTotalResource function is described as “Adjust overcommit metrics... The amount to overcommit will be re-calculated on next node heartbeat.”

wherein the supervisor controller is configured to monitor a second electronic random access memory capacity of a second computing system node,

The supervisor controller is configured to monitor a second electronic random access memory capacity of a second computer system node.

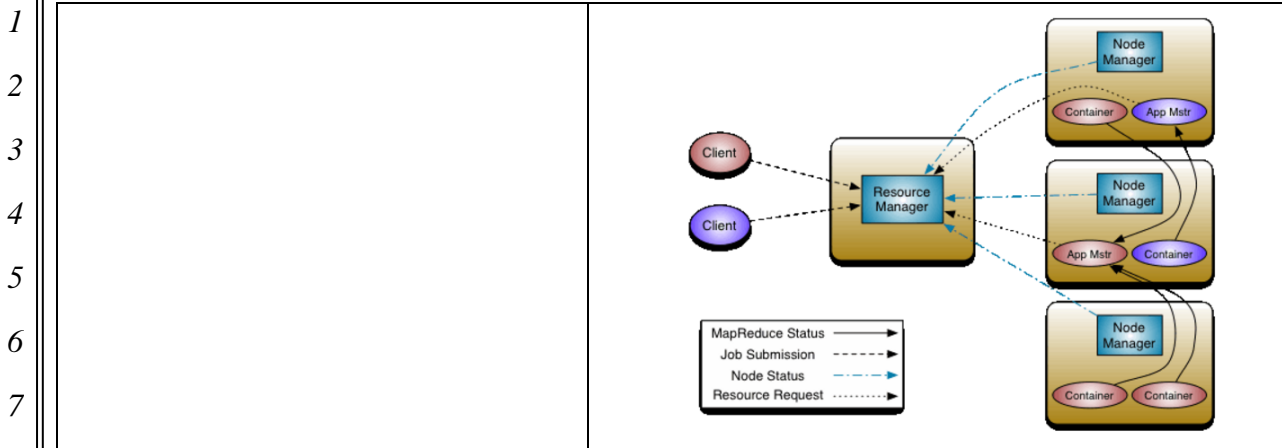
For example, the Hadoop 2.7.2 YARN Architecture<sup>5</sup> illustrates that the architecture includes multiple nodes:





1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

<p>wherein the supervisor controller is configured to prevent assignment of additional sub-jobs to a second computing system node based on determining that utilization of the second electronic random access memory capacity is at or above a threshold value, and</p>	<p>The supervisor controller is configured to prevent assignment of additional sub-jobs to a second computing system node based on determining that utilization of the second electronic random access memory capacity is at or above a threshold value.</p> <p>For example, in Yahoo’s dynamic overcommit patch, Yahoo’s source code comments explain that <code>RM_OVERCOMMIT_MEM_HWM</code> corresponds to: “High water memory utilization mark for overcommit. If the node's memory utilization is greater than or equal to this value then the node's memory will no longer be increased to allow further overcommit. A value <math>\leq 0</math> will disable memory overcommit.”</p> <p>Yahoo explained in its dynamic overcommit presentation: “If it's above the high watermark, then we're aggressively scaling back, meaning we don't want to allocate any more containers on that node at this point.” (See at 12 minute mark.)</p>
<p>wherein the master node and the plurality of computing system nodes include a computer processor and an electronic storage medium.</p>	<p>The master node and the plurality of computing system nodes include a computer processor and electronic storage medium.</p> <p>For example, the Hadoop 2.7.2 YARN Architecture illustrates:</p>



9           35.     However, despite being aware at least as early as August 4, 2016, of the '891  
10 and '593 patents and also aware of how its YARN-5202 patch falls within the scope of  
11 representative claims of both of those patents, Yahoo continued to use the YARN-5202 patch  
12 on Hadoop clusters in the United States at least until December 16, 2016, when Pepperdata  
13 commenced this action.

14           36.     Upon information and belief, Yahoo continues to use its YARN-5202 patch.

15           37.     Moreover, by presenting and explaining to numerous attendees at the 2016  
16 Hadoop Summit that using Yahoo's YARN-5202 patch provides significant improvements  
17 including increasing computing capacity in a Hadoop cluster by 50%, by publishing a video  
18 of that same presentation on the popular website youtube.com where it can be viewed for free  
19 by any member of the public with Internet access, and by making the YARN-5202 patch  
20 available for free Internet download to any member of the public by way of its JIRA, Yahoo  
21 actively encouraged Hadoop Summit attendees as well as other members of the public—  
22 particularly those operating Hadoop computer clusters—to download Yahoo's YARN-5202  
23 patch and use it in combination with a Hadoop computer cluster—and Yahoo did so while  
24 having knowledge of Pepperdata's '891 and '593 patents and while knowing that the use of  
25 its YARN-5202 patch in a Hadoop computer cluster falls within the scope of claims of  
26 Pepperdata's '891 and '593 patents.

27           38.     Upon and information and belief, as a result of Yahoo's active encouragement  
28 of Hadoop Summit attendees and others to download Yahoo's YARN-5202 patch and use it

1 in combination with a Hadoop computer cluster, Hadoop Summit attendees and others have  
2 in fact done so.

3 39. Yahoo knows that the YARN-5202 patch was created specifically to be used  
4 in combination with a Hadoop computer cluster and for the specific purpose of monitoring  
5 and allocating resources to facilitate overcommit functionality in a Hadoop computer cluster,  
6 and Yahoo knows that using the YARN-5202 patch in a Hadoop computer cluster to monitor  
7 and allocate resources to facilitate overcommit functionality falls within the scope of claims  
8 in Pepperdata's '891 and '593 patents and constitutes a material part of those claimed  
9 inventions. While having that knowledge, Yahoo offered the YARN-5202 patch for free use  
10 by others, making it available for free Internet download to any member of the public, and  
11 intending that it be combined with and used in a Hadoop computer cluster. Moreover, and  
12 again while having that same knowledge, Yahoo offered the YARN-5202 patch for free use  
13 by others outside of the United States, making it available for free Internet download to  
14 persons outside of the United States, and intending that it be combined with and used in a  
15 Hadoop computer cluster outside of the United States.

16 40. YARN-5202 is not a staple article or commodity of commerce, nor is it one  
17 that is suitable for substantial noninfringing use.

18 41. On information and belief, Yahoo has made, used, and made publicly available  
19 the YARN-5202 patch within and outside of the United States.

20 42. Weeks after Pepperdata's written notice to Yahoo of infringement of the '891  
21 and '593 patents, Yahoo informed Pepperdata that it had stopped making the YARN-5202  
22 patch available for download.

23  
24 **FIRST CLAIM FOR RELIEF:**

25 **INFRINGEMENT OF U.S. PATENT NO. 8,849,891**

26 43. Pepperdata incorporates paragraphs 1-42 of this Complaint.

27 44. This is a claim for patent infringement arising under the patent laws of the  
28 United States, Title 35 of the United States Code.

1           45. Without authority, Yahoo, through its agents, employees, and servants, has  
2 manufactured, used, promoted, offered for sale, and/or sold within the United States, and/or  
3 supplied in or from the United States, products and/or components covered by one or more  
4 claims of the '891 patent, and has, with knowledge of the '891 patent, actively induced others  
5 to do the same while knowing that the induced acts constituted infringement of the '891  
6 patent. Moreover, with knowledge of the '891 patent, Yahoo has provided products and  
7 components knowing that they, alone or as material components in combination with other  
8 components, infringe the '891 patent and has thereby contributed to others' infringement of  
9 the '891 patent. Yahoo has thereby infringed, actively induced others to infringe, and/or  
10 contributed to others' infringement of one or more claims of the '891 patent, including, for  
11 example and without limitation, claims 1-4 of the '891 patent, in violation of 35 U.S.C. § 271,  
12 including 35 U.S.C. §§ 271(a), (b), (c), and/or (f). This infringement is currently ongoing.  
13 The devices relating to Yahoo's infringement include, without limitation, Yahoo's Hadoop  
14 clusters using dynamic overcommit functionality, including those Hadoop clusters using the  
15 YARN-5202 software patch.

16           46. By no later than August 4, 2016, Pepperdata had given Yahoo written notice  
17 of its infringement of the '891 patent.

18           47. Yahoo's infringement of the '891 patent has been and continues to be  
19 deliberate and willful.

20           48. Yahoo's infringement of the '891 patent will continue unless enjoined by this  
21 Court.

22           49. Yahoo has derived and received, and will continue to derive and receive,  
23 gains, profits, and advantages from the aforesaid acts of infringement in an amount that is not  
24 presently known to Pepperdata.

25           50. Pepperdata has lost profits from the aforesaid acts of infringement in an  
26 amount that is not presently known to Pepperdata.

27           51. Due to Yahoo's infringement of the '891 patent, Pepperdata has been damaged  
28 and is entitled to monetary relief in an amount to be determined at trial.

1 52. Unless Yahoo is enjoined from infringing the '891 patent, Pepperdata will  
2 continue to suffer irreparable injury for which it has no adequate remedy at law.

3 **SECOND CLAIM FOR RELIEF:**

4 **INFRINGEMENT OF U.S. PATENT NO. 9,325,593**

5 53. Pepperdata incorporates paragraphs 1-42 of this Complaint.

6 54. This is a claim for patent infringement arising under the patent laws of the  
7 United States, Title 35 of the United States Code.

8 55. Without authority, Yahoo, through its agents, employees, and servants, has  
9 manufactured, used, promoted, offered for sale, and/or sold within the United States, and/or  
10 supplied in or from the United States, products and/or components covered by one or more  
11 claims of the '593 patent, and has, with knowledge of the '593 patent, actively induced others  
12 to do the same while knowing that the induced acts constituted infringement of the '593  
13 patent. Moreover, with knowledge of the '593 patent, Yahoo has provided products and  
14 components knowing that they, alone or as material components in combination with other  
15 components, infringe the '593 patent and has thereby contributed to others' infringement of  
16 the '593 patent. Yahoo has thereby infringed, actively induced others to infringe, and/or  
17 contributed to others' infringement of one or more claims of the '593 patent, including, for  
18 example and without limitation, claims 1-3 of the '593 patent, in violation of 35 U.S.C. § 271,  
19 including 35 U.S.C. §§ 271(a), (b), (c), and/or (f). This infringement is currently ongoing.  
20 The devices relating to Yahoo's infringement include, without limitation, Yahoo's Hadoop  
21 clusters using dynamic overcommit functionality, including those Hadoop clusters using the  
22 YARN-5202 software patch.

23 56. By no later than August 4, 2016, Pepperdata had given Yahoo written notice  
24 of its infringement of the '593 patent.

25 57. Yahoo's infringement of the '593 patent has been and continues to be  
26 deliberate and willful.

27 58. Yahoo's infringement of the '593 patent will continue unless enjoined by this  
28 Court.

1 59. Yahoo has derived and received, and will continue to derive and receive,  
2 gains, profits, and advantages from the aforesaid acts of infringement in an amount that is not  
3 presently known to Pepperdata.

4 60. Pepperdata has lost profits from the aforesaid acts of infringement in an  
5 amount that is not presently known to Pepperdata.

6 61. Due to Yahoo's infringement of the '593 patent, Pepperdata has been damaged  
7 and is entitled to monetary relief in an amount to be determined at trial.

8 62. Unless Yahoo is enjoined from infringing the '593 patent, Pepperdata will  
9 continue to suffer irreparable injury for which it has no adequate remedy at law.

10 **DEMAND FOR JUDGMENT**

11 Pepperdata respectfully prays for the following relief:

12 A. an order adjudging Yahoo to have infringed each of the '891 and '593 patents;

13 B. a permanent injunction enjoining Yahoo, as well as its officers, agents,  
14 servants, employees, and attorneys and those persons in active concert or participation with  
15 Yahoo, from infringing either or both of the '891 and '593 patents;

16 C. an accounting of all gains, profits, and advantages derived by Yahoo's  
17 infringement of the '891 and '593 patents and an award of damages adequate to compensate  
18 Pepperdata for that infringement;

19 D. an order trebling damages and/or for exemplary damages due to Yahoo's  
20 intentional and willful conduct;

21 E. an award of prejudgment and post judgment interest and costs to this action  
22 against Yahoo;

23 F. an award to Pepperdata of its attorneys' fees incurred in connection with this  
24 action; and

25 G. such other and further relief that the Court deems just and proper.

26 ///

27 ///

28 ///

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: February 10, 2017

By: /s/ Douglas G. Muehlhauser

Douglas G. Muehlhauser (State Bar No. 179,495)

doug.muehlhauser@knobbe.com

Alan G. Laquer (State Bar No. 259,257)

alan.laquer@knobbe.com

KNOBBE, MARTENS, OLSON & BEAR, LLP

2040 Main Street, Fourteenth Floor

Irvine, CA 92614

Telephone: 949-760-0404

Facsimile: 949-760-9502

*Attorneys for Plaintiff Pepperdata, Inc.*

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28

**JURY DEMAND**

Pursuant to Fed. R. Civ. P. 38(b), Plaintiff Pepperdata, Inc. demands a trial by jury of all issues raised by this Complaint that are triable by jury.

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: February 10, 2017

By: /s/ Douglas G. Muehlhauser  
Douglas G. Muehlhauser (State Bar No. 179,495)  
doug.muehlhauser@knobbe.com  
Alan G. Laquer (State Bar No. 259,257)  
alan.laquer@knobbe.com  
KNOBBE, MARTENS, OLSON & BEAR, LLP  
2040 Main Street, Fourteenth Floor  
Irvine, CA 92614  
Telephone: 949-760-0404  
Facsimile: 949-760-9502

*Attorneys for Plaintiff Pepperdata, Inc.*

25206186