

**UNITED STATES DISTRICT COURT  
FOR THE NORTHERN DISTRICT OF ILLINOIS  
EASTERN DIVISION**

Kove IO, Inc.,	)	
	)	
Plaintiff,	)	Civil Action No. 1:18-cv-08175
	)	
v.	)	Hon. Rebecca R. Pallmeyer
	)	
Amazon Web Services, Inc.,	)	
	)	
Defendant.	)	<b>DEMAND FOR JURY TRIAL</b>
	)	

**DEFENDANT AND COUNTERCLAIMANT AMAZON WEB SERVICES, INC.’S  
ANSWER, ADDITIONAL DEFENSES, AND COUNTERCLAIMS**

Defendant Amazon Web Services, Inc. (“AWS”), for its Answer, Additional Defenses, and Counterclaims to the Complaint dated December 12, 2018 (Dkt. 1), filed by Plaintiff Kove IO, Inc. (“Kove”), states as follows:

**NATURE OF THE ACTION**

1. Kove’s inventors developed breakthrough technology enabling high-performance, hyper-scalable distributed “cloud” storage years before the advent of the cloud. They were awarded patents for their innovations. Five years after this breakthrough, the *MIT Technology Review* recognized that the type of distributed data storage technology described in Kove’s patents was one of the top 10 emerging technologies that would change the world.<sup>1</sup>

**RESPONSE:** Denied.

2. Kove’s technology became essential to AWS as the volume of data stored on its cloud grew exponentially and its cloud storage business faced limitations on the ability to store and retrieve massive amounts of data. An AWS engineer explained that “we have a history of 10 years of using distributed hash tables to run systems and that becomes very powerful and that algorithm makes things work at massive scales and pretty much nothing else does.”<sup>2</sup> Once able to remove the barriers to massive data storage and retrieval present in previous systems, AWS became the first large-scale vendor of economical cloud infrastructure and services, allowing businesses

---

<sup>1</sup> **Plaintiff’s FN#1:** Balakrishnan, H., Distributed Storage, 10 Emerging Technologies That Will Change Your World, MIT Technology Review February 2004 (Attached as Ex. 1).

<sup>2</sup> **Plaintiff’s FN#2:** AWS re:Invent 2017: How DynamoDB Powered Amazon Prime Day 2017 (DAT326), YouTube (November 30, 2017), available at [https://www.youtube.com/watch?v=83-IWlvJ\\_8&feature=youtu.be&t=38m38s](https://www.youtube.com/watch?v=83-IWlvJ_8&feature=youtu.be&t=38m38s).

and individuals around the world access to the cloud without having to set up their own servers, software, and functionality. The ability to offer cloud services on this scope and scale was made possible through infringement of Kove's patents, paving the way for AWS to become what is believed to be Amazon's largest profit center.<sup>3</sup>

**RESPONSE:** To the extent Paragraph 2 purports to quote documents cited in the footnote to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Admitted that AWS is a cloud service vendor. Otherwise denied.

3. Kove is a Chicago-based pioneer in high-performance computer storage and data management technologies. It sells hardware and data management technology, creating solutions that overcome technological limitations in modern server architecture. It is a small, innovative product company competing in a field of behemoths, like AWS. It cannot outspend these dominant global players—respect for its intellectual property, as the law requires, is essential to fair competition. Companies such as AWS have little incentive to do business with small companies that have patented (and therefore disclosed) technology if they are able to take it without meaningful consequences. The disclosure of innovation in patents is not intended to facilitate unauthorized use, but rather to incentivize public disclosure for the benefit of all, in return for the promise to inventors of exclusive rights for a limited period of time.

**RESPONSE:** Denied.

4. Through this case, Kove seeks to stop AWS's unauthorized use of Kove's patented technology, and its noncompliance with the patent laws. This case is about ensuring a level playing field so smaller competitors like Kove can compete fairly on the basis of their hard work and protected innovations.

**RESPONSE:** Denied.

### **THE PARTIES**

5. Kove was founded in 2004 with an emphasis on high performance storage solutions. Kove has maintained a lean team of no more than 50 employees over the course of its history. The Kove team has included those with strong backgrounds from premier tech companies, including Intel, Sun, Cisco, Seagate, Quantum, NASA, and many others. Kove's headquarters are located in Chicago, Illinois. It owns the patents asserted in this case. Dr. John Overton and Dr. Stephen Bailey are the named inventors, and Dr. Overton serves as Kove's Chief Executive Officer.

---

<sup>3</sup> **Plaintiff's FN#3:** Wingfield, N., Amazon's Profit Swells to \$1.6 Billion, Lifted by Its Cloud Business, New York Times, (April 26, 2018), available at <https://www.nytimes.com/2018/04/26/technology/amazon-prime-profit.html> (Attached as Ex. 2).

**RESPONSE:** Admitted that John K. Overton and Stephen W. Bailey are the listed inventors on U.S. Patent Nos. 7,103,640; 7,233,978; and 7,814,170. Otherwise, AWS is without knowledge or information sufficient to form a belief as to the truth of the allegations in this paragraph, and therefore denies those allegations.

6. Kove is a corporation organized and existing under the laws of the State of Delaware and is registered to do business in the State of Illinois. Kove's principal place of business is located at 14 North Peoria Street Suite 2H, Chicago, Illinois 60607.

**RESPONSE:** AWS is without knowledge or information sufficient to form a belief as to the truth of the allegations in this paragraph, and therefore denies those allegations.

7. AWS is a corporation organized and existing under the laws of the State of Delaware and is registered to do business in the State of Illinois. Upon information and belief, AWS has its principal place of business at 410 Terry Avenue North Seattle, Washington 98109 and a regular and established place of business in this District at AT&T Center, 227 W. Monroe Street, Chicago, Illinois 60606.

**RESPONSE:** Admitted.

8. AWS develops and provides cloud storage products and services to customers, including customers in this District.

**RESPONSE:** Admitted.

### **JURISDICTION AND VENUE**

9. This Court has subject matter jurisdiction over this case under 28 U.S.C. §§ 1331 and 1338. Venue is proper in this Court under 28 U.S.C. §§ 1391 and 1400(b).

**RESPONSE:** Admitted that the Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338. Otherwise, denied.

10. This Court has personal jurisdiction over AWS. AWS has continuous and systematic business contacts with the State of Illinois. AWS, directly and/or through subsidiaries or intermediaries, conducts its business extensively throughout Illinois, by shipping, distributing, offering for sale, selling, and advertising (including the provision of interactive web pages) its products and services in the State of Illinois and in this District. AWS, directly and/or through subsidiaries or intermediaries, has purposefully and voluntarily placed its infringing products and services into this District and into the stream of commerce with the intention and expectation that they will be purchased and used by consumers in this District. AWS has offered and sold and

continues to offer and sell these infringing products and services in this District.<sup>4</sup> On information and belief, AWS, for example, sells and offers to sell the infringing products and services to developers, partners or, customers in this District, such as Amazon.com, Inc., University of Chicago, City of Chicago, and Federal Home Loan Bank of Chicago. AWS has committed acts of infringement in this District and has a regular and established place of business in this District.

**RESPONSE:** Denied.

### **BACKGROUND**

11. Drs. Overton and Bailey, who met at the University of Chicago while working on their PhDs, are the named inventors of the three patents-in-suit: U.S. Patent Nos. 7,814,170 (“‘170 Patent”); 7,103,640 (“‘640 Patent”); and 7,233,978 (“‘978 Patent”) (collectively, “patents-in-suit”).

**RESPONSE:** Admitted that John K. Overton and Stephen W. Bailey are the listed inventors on U.S. Patent Nos. 7,103,640; 7,233,978; and 7,814,170. Otherwise, AWS is without knowledge or information sufficient to form a belief as to the truth of the allegations in this paragraph, and therefore denies those allegations.

12. Today, cloud storage is so ubiquitous that it has become part of the common vernacular. When the inventors were studying together at the University of Chicago in the 1990’s, the cloud as we know it was years away. The inventors foresaw the advent of the cloud—distributed, large-scale storage networks—and they identified a key roadblock that would need to be overcome in order to sustain viability.

**RESPONSE:** Denied.

13. In particular, data storage management has become increasingly critical for companies as the amount of data produced every day grows exponentially—for example, today, there are 2.5 quintillion (2.5 \*10<sup>18</sup>) bytes of data created each day.<sup>5</sup> In the 1990’s, the inventors foresaw that data storage requirements would grow beyond the capabilities of conventional computer networks. Dependency on traditional ways of indexing the information in a distributed storage network was identified as a key issue. For example, as information was added to a network, previous technology would require that it be indexed in some hierarchical fashion, and those indices would periodically have to be updated. While these indices worked well to a certain point,

---

<sup>4</sup> **Plaintiff’s FN#4:** AWS, Customer Success Stories, available at <https://aws.amazon.com/solutions/case-studies/startups/> and <https://aws.amazon.com/solutions/case-studies/government-education/all-government-education-nonprofit/> (Attached as Ex. 3).

<sup>5</sup> **Plaintiff’s FN#5:** Bernard Marr, How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read (May 21, 2018), available at <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#2329790360ba>.

the distributed network of data servers and data would become so vast and complex that ordinary hierarchical indices no longer could handle the load without high levels of inefficiency.

**RESPONSE:** Denied.

14. Further, to store a piece of information (e.g., a data file), a storage system must store not only the data file itself but also its corresponding location information, which records where the data file is located on the network of servers and computers. Without the corresponding location information, a storage system would not know where to find the data file (i.e., which server on the network to access and from which to retrieve the data file).

**RESPONSE:** Denied.

15. Location information of data files was traditionally stored on a single centralized server, and it was retrieved from that specific centralized server. Thus, then-existing systems allowed a client seeking location information associated with a data file to send a request to a server, and typically, only data statically associated with that server was returned. The search was conducted only where the system knew in advance to look.

**RESPONSE:** Denied.

16. In the 1990's, such a server may have been sufficient; however, the inventors knew distributed storage systems would someday contain so many unique data files that it would become impractical—if not impossible—to store the corresponding location information in one place.

**RESPONSE:** Denied.

17. Drs. Overton and Bailey addressed these challenges. They realized, among other things, that storing location information associated with data files across multiple servers would reduce the processing time to find a data file. Likewise, the inventors understood the need to efficiently identify which of the multiple location information servers stored the location information for a particular data file.

**RESPONSE:** Denied.

18. To achieve this, Drs. Overton and Bailey invented the claimed technology of the patents-in-suit. For example, hash values corresponding to the location information of data files would be stored and distributed on the network. The distributed hash values would reveal which of the location information servers contained the location information of particular data files. Each hash value may point to the location information server storing the location information of a data file. Each location server would have the ability to point to other servers based on the hash values. While the first server may not have the location information for the file the user is looking for, it will reroute—within milliseconds—to another server that does have that location information. This fundamental technology allowed for the efficient organization of location information even as files are constantly added, deleted, modified, and moved. It enabled hyper-scalable cloud storage and improved upon the scalability limitations of conventional storage systems.

**RESPONSE:** Denied.

19. The claimed inventions were ahead of their time—distributed storage networks had not yet grown large enough. Years after the inventors filed the earliest of the original patent applications, an article published in the *MIT Technology Review* described the use of distributed hash tables (as discussed above) as one of “10 Emerging Technologies That Will Change Your World,” explaining that “[h]ash tables provide a quick way to organize data: a simple mathematical operation assigns each file its own row in a table; the row stores the file’s location,” and that “distributed hash tables [are] the coming future of networked storage” and have a “huge variety of applications. . . . Not very many technologies have such broad potential.”<sup>6</sup>

**RESPONSE:** To the extent Paragraph 19 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS denies that this Paragraph accurately quotes the cited excerpts, and refers the Court to those documents for their full content. Otherwise denied.

### **THE KOVE PATENTS-IN-SUIT**

20. On September 5, 2006, the U.S. Patent and Trademark Office duly and legally issued the ’640 patent, entitled “Network Distributed Tracking Wire Transfer Protocol,” with Drs. Overton and Bailey as inventors. The earliest application related to the ’640 patent was filed on July 8, 1998. A true and correct copy of the ’640 patent is attached as Exhibit 4.

**RESPONSE:** Admitted that the ’640 patent was issued on September 5, 2006, is entitled “Network Distributed Tracking Wire Transfer Protocol,” and lists John K. Overton and Stephen W. Bailey as inventors. Otherwise denied.

21. On June 19, 2007, the U.S. Patent and Trademark Office duly and legally issued the ’978 patent, entitled “Method and Apparatus for Managing Location Information in a Network Separate From the Data to Which the Location Information Pertains,” with Drs. Overton and Bailey as inventors. The earliest application related to the ’978 patent was filed on July 8, 1998. A true and correct copy of the ’978 patent is attached as Exhibit 5.

**RESPONSE:** Admitted that the ’978 patent was issued on June 19, 2007, is entitled “Method and Apparatus for Managing Location Information in a Network Separate From the Data to Which

---

<sup>6</sup> **Plaintiff’s FN#6:** *See supra* n. 1. (Balakrishnan, Ex. 1).

the Location Information Pertains,” and lists John K. Overton and Stephen W. Bailey as inventors.

Otherwise denied.

22. On October 12, 2010, the U.S. Patent and Trademark Office duly and legally issued the '170 patent, entitled “Network Distributed Tracking Wire Transfer Protocol,” with Drs. Overton and Bailey as inventors. The earliest application related to the '170 patent was filed on July 8, 1998. A true and correct copy of the '170 patent is attached as Exhibit 6.

**RESPONSE:** Admitted that the '170 patent was issued on October 12, 2010, is entitled “Network Distributed Tracking Wire Transfer Protocol,” and lists John K. Overton and Stephen W. Bailey as inventors. Otherwise denied.

23. Kove is the sole and exclusive owner of all rights, title, and interest to the patents-in-suit necessary to bring this action, including the right to recover past and future damages.

**RESPONSE:** AWS is without knowledge or information sufficient to form a belief as to the truth of the allegations in this paragraph, and therefore denies those allegations.

24. The patents-in-suit are enforceable and valid.

**RESPONSE:** Denied.

#### **AWS'S INFRINGING PRODUCTS AND SERVICES**

25. Upon information and belief, AWS has infringed and continues to infringe, directly and indirectly, one or more claims of the patents-in-suit.

**RESPONSE:** Denied.

26. The accused products include without limitation cloud-based products and services provided by AWS, such as Amazon Simple Storage Service (“Amazon S3”), DynamoDB, and other related products and services (collectively, the “AWS Accused Products”).

**RESPONSE:** Admitted that Plaintiff accuses Amazon S3 and DynamoDB of infringement.

Otherwise denied.

27. As an example, Amazon S3 allows users to store data in the cloud. In fact, it is described as “storage for the internet.” Specifically, it launched in 2006 and was designed to provide highly scalable, reliable, and low-latency data storage infrastructure at very low costs, where users are able to “[w]rite, read, and delete objects containing from 1 byte to 5 gigabytes of



data each” and where “[e]ach object is stored and retrieved via a unique developer-assigned key.”<sup>7</sup> In the first seven years of operation, Amazon S3 experienced a 20,000% increase in the number of data objects stored, reaching 2 trillion ( $2 \times 10^{12}$ ) data objects by 2013.<sup>8</sup> Amazon S3 is considered “one of the wonders” of AWS.<sup>9</sup>

**RESPONSE:** To the extent Paragraph 27 quotes from excerpts of documents cited in footnote 7 to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Denied that footnote 9 contains the quoted language. Otherwise denied.

28. As another example, DynamoDB is described as a “fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. DynamoDB lets you offload the administrative burdens of operating and scaling a distributed database, so that you don’t have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling. With DynamoDB, you can create database tables that can store and retrieve any amount of data, and serve any level of request traffic. You can scale up or scale down your tables’ throughput capacity without downtime or performance degradation, and use the AWS Management Console to monitor resource utilization and performance metrics.”<sup>10</sup>

**RESPONSE:** To the extent Paragraph 28 purports to quote documents cited in the footnote to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Otherwise denied.

29. “DynamoDB powers the next wave of high-performance, internet-scale applications. . . . DynamoDB is used by Lyft to store GPS locations for all their rides, Tinder to store millions of user profiles and make billions of matches, Redfin to scale to millions of users and manage data for hundreds of millions of properties, Comcast to power their XFINITY X1

---

<sup>7</sup> **Plaintiff’s FN#7:** Amazon Web Services Launches, Press Release (March 14, 2006), at \*1, available at <http://phx.corporate-ir.net/phoenix.zhtml?c=176060&p=irol-newsArticle&ID=830816> (Attached as Ex. 7); Feloni, R., How A Regular Employee Helped Put Amazon On The Path To Billions Of Dollars, Business Insider (Jul. 22, 2014), available at <http://www.businessinsider.com/benjamin-black-and-amazon-web-services-2014-7> (Attached as Ex. 8).

<sup>8</sup> **Plaintiff’s FN#8:** Barr, J., AWS News Blog, Amazon S3 – Two Trillion Objects, 1.1 Million Requests / Second, available at <https://aws.amazon.com/blogs/aws/amazon-s3-two-trillion-objects-1.1-million-requests-second/> (Attached as Ex. 9).

<sup>9</sup> **Plaintiff’s FN#9:** Hern, A., Amazon Web Services: the secret to the online retailer’s future success, The Guardian (February 2, 2017), available at <https://www.theguardian.com/technology/2017/feb/02/amazon-web-services-the-secret-to-the-online-retailers-future-success> (Attached as Ex. 10).

<sup>10</sup> **Plaintiff’s FN#10:** DynamoDB Developer Guide (API Version 2012-08-10), at \*1, available at <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/dynamodb-dg.pdf> (While the full document is incorporated by reference herein, an excerpted version of the document is attached as Ex. 11); see also Amazon’s Web Version of the DynamoDB Developer Guide at <https://docs.aws.amazon.com/amazon-dynamodb/latest/developerguide/Introduction.html>.



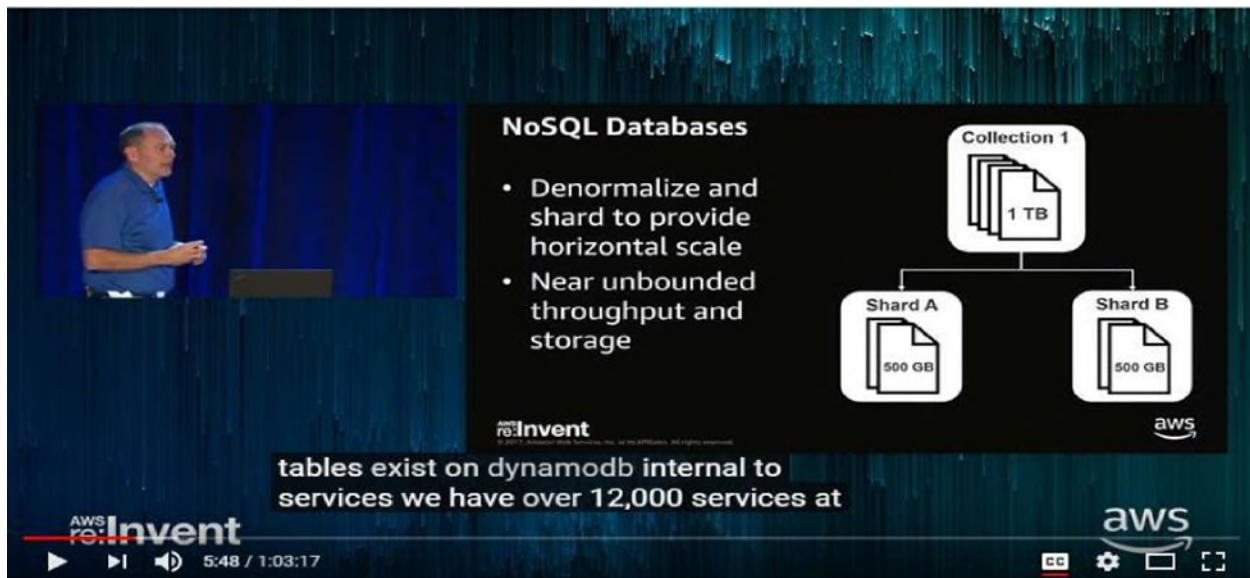
video service running on more than 20 million devices, BMW to run its car-as-a-sensor service that can scale up and down by two orders of magnitude within 24 hours, Nordstrom for their recommendations engine reducing processing time from 20 minutes to a few seconds, Under Armour to support its connected fitness community of 200 million users, Toyota Racing to make real time decisions on pit-stops, tire changes, and race strategy, and another 100,000+ AWS customers for a wide variety of high-scale, high-performance use cases.”<sup>11</sup>

**RESPONSE:** To the extent Paragraph 29 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

30. On information and belief, AWS uses the AWS Accused Products for its own business purposes. In addition, AWS regularly conducts testing and troubleshooting of the AWS Accused Products. Further, Kove is informed and believes companies related to AWS (e.g., Amazon.com, Inc. and its subsidiaries) use the AWS Accused Products.


**RESPONSE:** Denied.

31. On information and belief, the features discussed herein are not limited to any one of the AWS Accused Products. For example, as an AWS senior practice manager explained: “[T]ables exist on DynamoDB internal to services we have over 12,000 services at AWS many many of those services use DynamoDB . . . .”<sup>12</sup>



<sup>11</sup> **Plaintiff’s FN#11:** Vogels, W., A Decade of Dynamo: Powering the net wave of high-performance, internet-scale applications, All Things Distributed, at \*2, available at <https://www.allthingsdistributed.com/2017/10/a-decade-of-dynamo.html> (Attached as Ex. 12).

<sup>12</sup> **Plaintiff’s FN#12:** AWS Youtube Videos, AWS re:Invent 2017: DynamoDB adaptive capacity: smooth performance for chaotic workl (DAT327), available at [https://youtu.be/kMY0\\_m29YzU](https://youtu.be/kMY0_m29YzU).



**RESPONSE:** To the extent Paragraph 31 quotes from excerpts and reproduces screenshots of the video cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes and reproduces screenshots of the video cited, but refers the Court to that video for its full content. Otherwise denied.

32. Indeed, on information and belief, the AWS Accused Products are based on a common platform, like the “Dynamo” technology: “Dynamo is internal technology developed at Amazon to address the need for an incrementally scalable, highly-available key-value storage system. The technology is designed to give its users the ability to trade-off cost, consistency, durability and performance, while maintaining high-availability.”<sup>13</sup> Dynamo is used “to power parts of [AWS products], such as S3.”<sup>14</sup> DynamoDB too is based on Dynamo technology and is built on AWS products such as S3.<sup>15</sup>

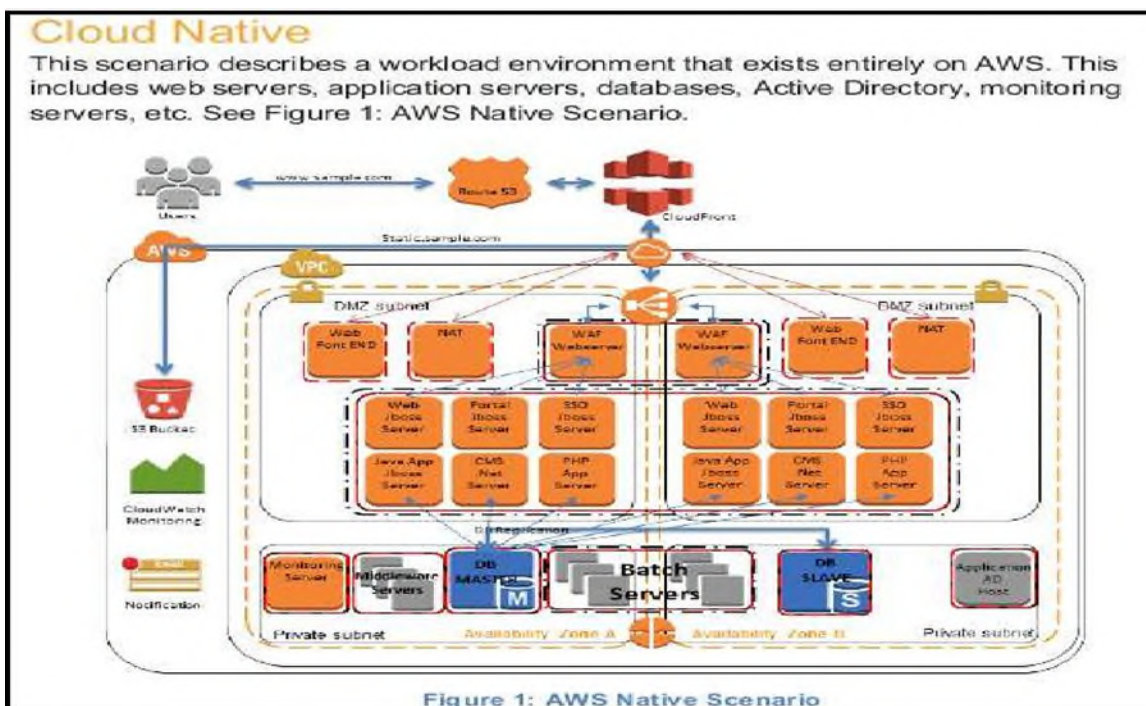
<sup>13</sup> **Plaintiff’s FN#13:** Vogels, W., Amazon’s Dynamo, All Things Distributed, at \*1, available at [https://www.allthingsdistributed.com/2007/10/amazons\\_dynamo.html](https://www.allthingsdistributed.com/2007/10/amazons_dynamo.html) (Attached as Ex. 13).

<sup>14</sup> **Plaintiff’s FN#14:** *Id.*

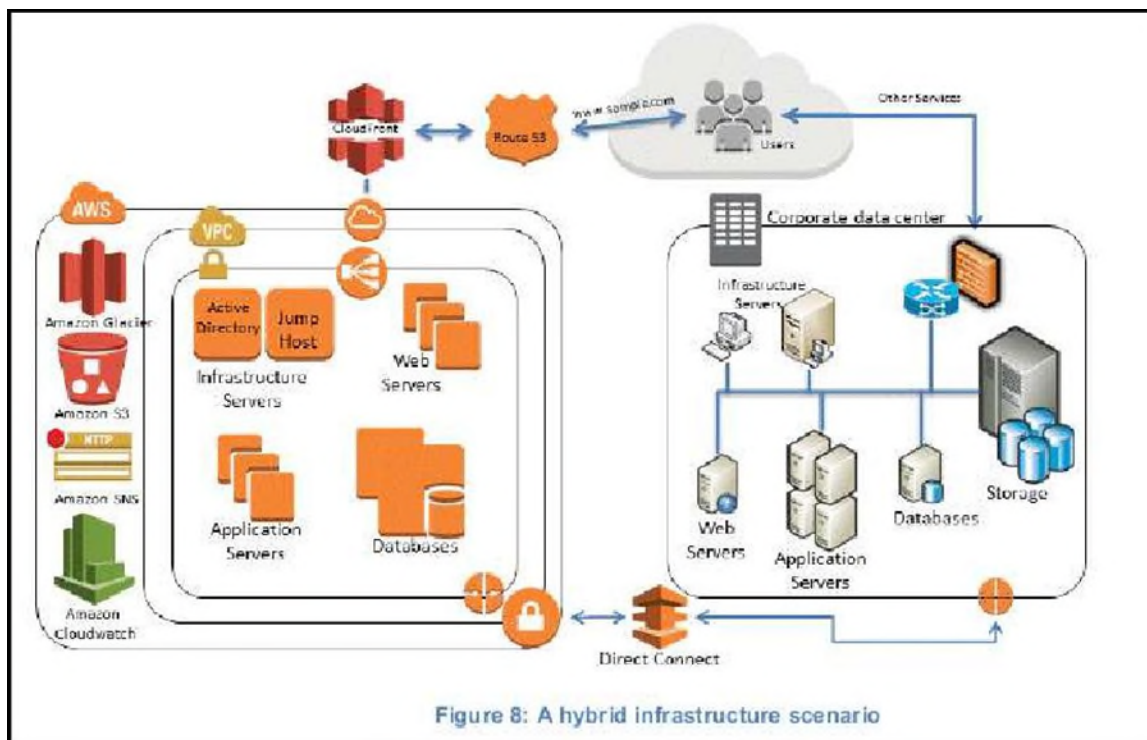
<sup>15</sup> **Plaintiff’s FN#15:** Vogels, W., Amazon DynamoDB – a Fast and Scalable NoSQL Database Service Designed for Internet Scale Applications, All Things Distributed, at \*1, available at <https://www.allthingsdistributed.com/2012/01/amazon-dynamodb.html> (Attached as Ex. 14).

**RESPONSE:** To the extent Paragraph 32 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

33. On information and belief, AWS offers an easily integrated platform, consisting of the AWS Accused Products, that uses the features discussed herein as shown below:<sup>16</sup>



<sup>16</sup> **Plaintiff's FN#16:** Agnihotri, P., Amazon Web Services – Backup, Archive and Restore Approaches Using AWS, at \*5, \*16 available at [https://d0.awsstatic.com/whitepapers/Backup\\_Archive\\_and\\_Restore\\_Approaches\\_Using\\_AWS.pdf](https://d0.awsstatic.com/whitepapers/Backup_Archive_and_Restore_Approaches_Using_AWS.pdf) (Attached as Ex. 15).



**RESPONSE:** To the extent Paragraph 33 reproduces figures in documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately reproduces the cited figures, but refers the Court to those documents for their full content. Otherwise denied.

### COUNT I

#### AWS’S INFRINGEMENT OF THE ’170 PATENT

34. Kove incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

**RESPONSE:** AWS incorporates herein by reference its responses to Paragraphs 1 through 33.

35. On information and belief, AWS has infringed and will continue to infringe the ’170 patent. AWS directly infringes the ’170 patent under 35 U.S.C. §271(a) by making, using, selling, offering for sale, and/or importing in this District and into the United States products and/or methods covered by one or more claims of the ’170 patent, including, but not limited to, the AWS Accused Products. As an example, the AWS Accused Products infringe at least claim 1 of the ’170 patent.

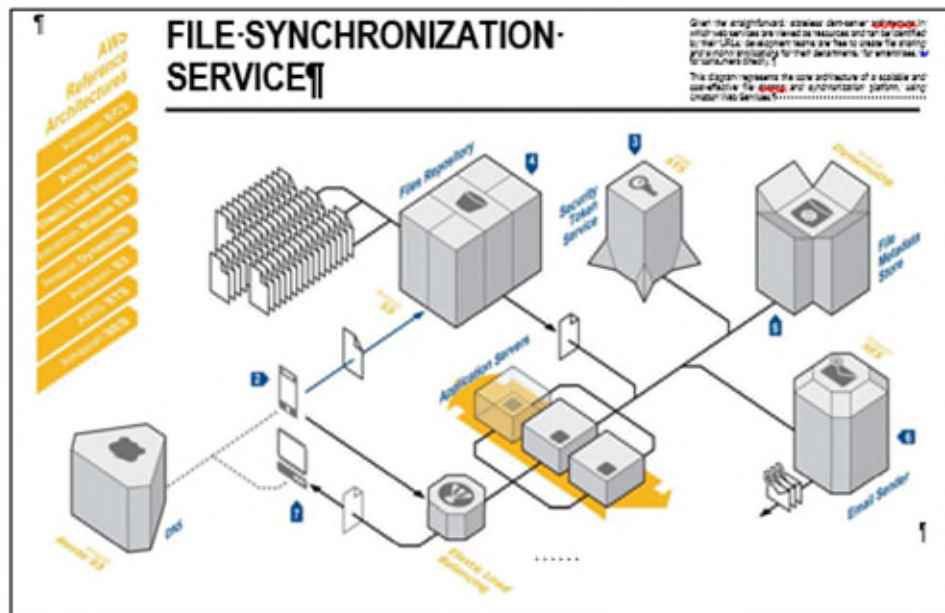


**RESPONSE:** Denied.

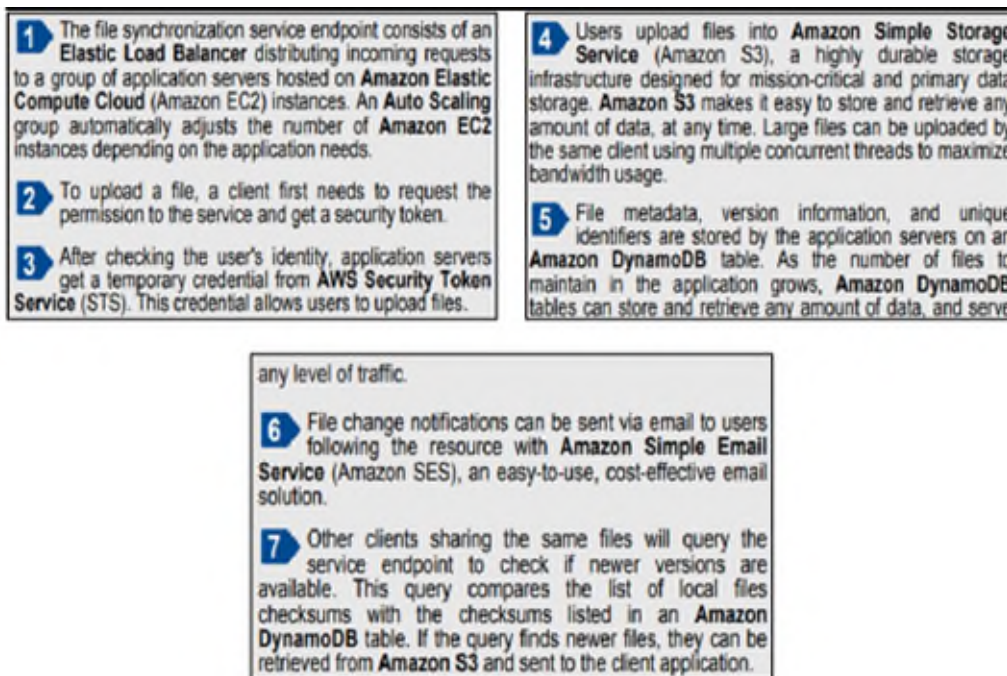
36. Claim 1 is directed to a “system for managing data stored in a distributed network.” Claim 1, for example, recites a “data repository configured to store a data entity, wherein an identifier string identifies the data entity.” Claim 1 further recites a “data location server network” where “data location information for a plurality of data entities is stored.” Further still, claim 1 recites a “hash function used to organize the data location information across the . . . data location servers.”

**RESPONSE:** To the extent Paragraph 36 quotes from excerpts of Claim 1 of the ’170 patent, AWS admits that this Paragraph accurately quotes the excerpts, but refers the Court to the ’170 patent for its full content. Otherwise denied.

37. The AWS Accused Products manage data stored in a distributed network. For example, as shown below in the Amazon S3 Reference Architecture diagram (“Architecture Diagram”), “[u]sers upload files into Amazon Simple Storage Service (Amazon S3), a highly durable storage infrastructure designed for mission-critical and primary data storage. Amazon S3 makes it easy to store and retrieve any amount of data, at any time.”<sup>17</sup>



<sup>17</sup> **Plaintiff’s FN#17:** Amazon Web Services, Inc., [https://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_filesync\\_08.pdf](https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf) (modified from original) Attached as Ex. 16).



**RESPONSE:** To the extent Paragraph 37 quotes from excerpts and reproduces figures of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts and reproduces the cited figures, but refers the Court to those documents for their full content. Otherwise denied.

38. AWS stores data entities, wherein an identifier string identifies the data entities. In particular, Amazon S3 stores objects in a “Files Repository,” and each Amazon S3 object, is “uniquely identified within a bucket by a key (name) and a version ID. . . . Amazon S3 can be thought of as a basic data map between ‘bucket + key + version’ and the object itself. Every object in Amazon S3 can be uniquely addressed through the combination of the web service endpoint, bucket name, key, and optionally, a version.”<sup>18</sup>

**RESPONSE:** To the extent Paragraph 38 purports to quote documents cited in the footnote to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Otherwise denied.

39. On information and belief, AWS includes a data location server network where data location information for a plurality of data entities is stored. For example, Amazon S3 maintains

<sup>18</sup> **Plaintiff’s FN#18:** Amazon Simple Storage Service Developer Guide, (API Version 2006-03-01), at \*3–4, available at <http://docs.aws.amazon.com/AmazonS3/latest/dev/s3-dg.pdf> (While the full document is incorporated by reference herein, an excerpted version of the document is attached as Ex. 17); *see also* Amazon’s Web Version of the S3 Developer Guide at <https://docs.aws.amazon.com/AmazonS3/latest/dev/Welcome.html>.

an index subsystem that manages the metadata and location information of all S3 objects in the data location servers as illustrated below.<sup>19</sup>

We'd like to give you some additional information about the service disruption that occurred in the Northern Virginia (US-EAST-1) Region on the morning of February 28th, 2017. The Amazon Simple Storage Service (S3) team was debugging an issue causing the S3 billing system to progress more slowly than expected. At 9:37AM PST, an authorized S3 team member using an established playbook executed a command which was intended to remove a small number of servers for one of the S3 subsystems that is used by the S3 billing process. Unfortunately, one of the inputs to the command was entered incorrectly and a larger set of servers was removed than intended. The servers that were inadvertently removed supported two other S3 subsystems: **One of these subsystems, the index subsystem, manages the metadata and location information of all S3 objects in the region.** This subsystem is necessary to serve all GET, LIST, PUT, and DELETE requests. The second subsystem, the placement subsystem, manages allocation of new storage and requires the index subsystem to be functioning properly to correctly operate. The placement subsystem is used during PUT requests to allocate storage for new objects. Removing a significant portion of the capacity caused each of these systems to require a full restart. While these subsystems were being restarted, S3 was unable to service requests. Other AWS services in the US-EAST-1 Region that rely on S3 for storage, including the S3 console, Amazon Elastic Compute Cloud (EC2) new instance launches, Amazon Elastic Block Store (EBS) volumes (when data was needed from a S3 snapshot), and AWS Lambda were also impacted while the S3 APIs were unavailable.

**RESPONSE:** To the extent Paragraph 39 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

40. On information and belief, Amazon S3 stores data location information for S3 objects in a DynamoDB table.<sup>20</sup> It has been stated that “[t]he heart of the S3 object index is a DynamoDB table with one item per object, which associates various attributes with the object’s S3 key. Each item contains the S3 key, the size of the object, and any additional attributes to use for lookups.”<sup>21</sup>

## Storing Large Attribute Values in Amazon S3

As mentioned above, you can also take advantage of Amazon Simple Storage Service (Amazon S3) to store large attribute values that cannot fit in a DynamoDB item. You can store them as an object in Amazon S3 and then store the object identifier in your DynamoDB item.

You can also use the object metadata support in Amazon S3 to provide a link back to the parent item in DynamoDB. Store the primary key value of the item as Amazon S3 metadata of the object in Amazon S3. Doing this often helps with maintenance of the Amazon S3 objects.

<sup>19</sup> **Plaintiff’s FN#19:** Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region, available at <https://aws.amazon.com/message/41926/> (emphasis added) (Attached as Ex. 18).

<sup>20</sup> **Plaintiff’s FN#20:** See *supra* n. 10 at \*523 (DynamoDB Developer Guide, Ex. 11).

<sup>21</sup> **Plaintiff’s FN#21:** Deck, M., AWS Big Data Blog, Building and Maintaining an Amazon S3 Metadata Index without Servers, available at <https://aws.amazon.com/blogs/big-data/building-and-maintaining-an-amazon-s3-metadata-index-without-servers/> (Attached as Ex. 19).



**RESPONSE:** To the extent Paragraph 40 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Denied that the screenshot excerpted is contained in the documents cited. Otherwise denied.

41. AWS’s “Frequently Asked Questions About Amazon DynamoDB” explains that “file pointers (possibly to Amazon S3 objects) are best saved in Amazon DynamoDB,” as opposed to in Amazon S3.<sup>22</sup> AWS allows “the object metadata support in Amazon S3 to store the primary key value of the corresponding item as Amazon S3 object metadata. This use of metadata can help with future maintenance of your Amazon S3 objects.”<sup>23</sup>

**RESPONSE:** To the extent Paragraph 41 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Denied that the webpages at the cited links contain the quoted language. Otherwise denied.

42. On information and belief, AWS uses a hash function to organize data location information across the data location servers. For example, DynamoDB tables organize data location information across servers based on a hash function. AWS has described DynamoDB as a system that “uses its hash function to determine where to store a new item . . . . Note that the items are not stored in sorted order. Each item’s location is determined by the hash value of its partition key.”<sup>24</sup>

---

<sup>22</sup> **Plaintiff’s FN#22:** Amazon DynamoDB FAQs, at \*4–5, available at <https://aws.amazon.com/dynamodb/faqs/> (Attached as Ex. 20).

<sup>23</sup> **Plaintiff’s FN#23:** See *supra* n. 10 at \*523 (DynamoDB Developer Guide, Ex. 11).

<sup>24</sup> **Plaintiff’s FN#24:** AWS Documentation, DynamoDB Core Components, available at <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.CoreComponents.html#HowItWorks.CoreComponents.PrimaryKey>; see also *supra* n. 10 at \*2–9.

**Primary Key**

When you create a table, in addition to the table name, you must specify the primary key of the table. The primary key uniquely identifies each item in the table, so that no two items can have the same key.

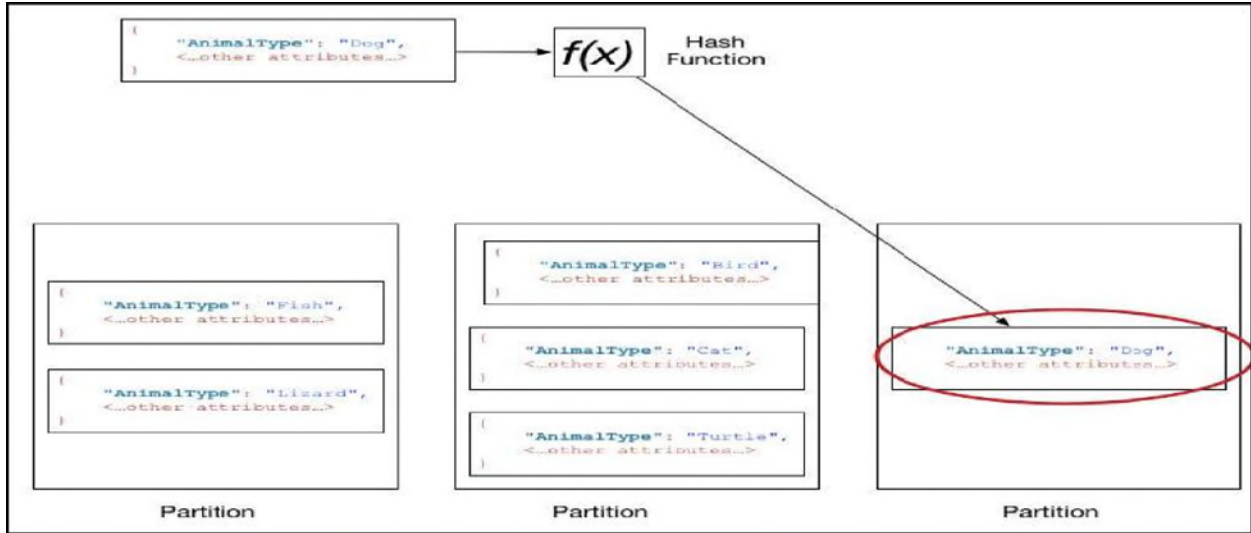
DynamoDB supports two different kinds of primary keys:

- **Partition key** – A simple primary key, composed of one attribute known as the *partition key*. DynamoDB uses the partition key's value as input to an internal hash function. The output from the hash function determines the partition (physical storage internal to DynamoDB) in which the item will be stored. In a table that has only a partition key, no two items can have the same partition key value. The *People* table described in [Tables, Items, and Attributes](#) is an example of a table with a simple primary key (*PersonID*). You can access any item in the *People* table directly by providing the *PersonID* value for that item.
- **Partition key and sort key** – Referred to as a *composite primary key*, this type of key is composed of two attributes. The first attribute is the *partition key*, and the second attribute is the *sort key*. DynamoDB uses the partition key value as input to an internal hash function. The output from the hash function determines the partition (physical storage internal to DynamoDB) in which the item will be stored. All items with the same partition key are stored together, in sorted order by sort key value. In a table that has a partition key and a sort key, it's possible for two items to have the same partition key value. However, those two items must have different sort key values. The *Music* table described in [Tables, Items, and Attributes](#) is an example of a table with a composite primary key (*Artist* and *SongTitle*). You can access any item in the *Music* table directly, if you provide the *Artist* and *SongTitle* values for that item. A composite primary key gives you additional flexibility when querying data. For example, if you provide only the value for *Artist*, DynamoDB retrieves all of the songs by that artist. To retrieve only a subset of songs by a particular artist, you can provide a value for *Artist* along with a range of values for *SongTitle*.

**RESPONSE:** To the extent Paragraph 42 quotes from excerpts of documents and reproduces figures cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts and reproduces the excerpted figure, but refers the Court to those documents for their full content. Otherwise denied.

43. As AWS illustrates in its DynamoDB Developer Guide, “[t]he following diagram [below] shows a table named PETS, which spans multiple partitions. The table’s primary key is ANIMALTYPE (only this key attribute is shown). DynamoDB uses its hash function to determine where to store a new item, in this case based on the hash value of the string DOG. Note that the items are not stored in sorted order. Each item’s location is determined by the hash value of its partition key.”<sup>25</sup>

<sup>25</sup> **Plaintiff’s FN#25:** See *supra* n. 10 at \*18–20.



**RESPONSE:** To the extent Paragraph 43 quotes from excerpts and reproduces figures of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts and reproduces the excerpted figure, but refers the Court to those documents for their full content. Otherwise denied.

44. “To read that same item from the PETS table, DynamoDB calculates the hash value of DOG, yielding the partition in which these items are stored.”<sup>26</sup> On information and belief, these partitions (e.g., nodes) reside on AWS servers.

**RESPONSE:** To the extent Paragraph 44 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

45. It has been further described that “DynamoDB automatically spreads the data and traffic for your tables over a sufficient number of servers to handle your throughput and storage requirements, while maintaining consistent and fast performance. All of your data is stored on solid state disks (SSDs) and automatically replicated across multiple Availability Zones in an AWS region, providing built-in high availability and data durability.”<sup>27</sup>

<sup>26</sup> **Plaintiff’s FN#26:** *Id.* at \*20.

<sup>27</sup> **Plaintiff’s FN#27:** *Id.* at \*1.

**RESPONSE:** To the extent Paragraph 45 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

46. Moreover, on information and belief, AWS receives read and store (i.e., “get” and “put”) operations for the keys, e.g., the ANIMALTYPE example discussed above: “Both get and put operations are invoked using Amazon’s infrastructure-specific request processing framework over HTTP. There are two strategies that a client can use to select a node: (1) route its request through a generic load balancer that will select a node based on load information, or (2) use a partition-aware client library that routes requests directly to the appropriate coordinator nodes. . . . A node handling a read or write operation is known as the coordinator. Typically, this is the first among the top N nodes in the preference list. If the requests are received through a load balancer, requests to access a key may be routed to any random node in the ring.”<sup>28</sup>

**RESPONSE:** To the extent Paragraph 46 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

47. Further still, it has been described that “partitioning and placement information also propagates via the gossip-based protocol and each storage node is aware of the token ranges handled by its peers. This allows each node to forward a key’s read/write operations to the right set of nodes directly.”<sup>29</sup>

**RESPONSE:** To the extent Paragraph 47 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

48. Based on the above, the AWS Accused Products directly infringe at least, but not limited to, claim 1 of the ’170 patent.

**RESPONSE:** Denied.

49. AWS also indirectly infringes the ’170 patent by inducing others to infringe and/or contributing to the infringement of others, including third party users of the AWS Accused Products in this District and elsewhere in the United States. Specifically, on information and belief,

---

<sup>28</sup> **Plaintiff’s FN#28:** DeCandia, G., et al., Dynamo: Amazon’s Highly Available Key-value Store, Amazon.com (Oct. 14–17, 2007) at \*211 (Attached as Ex. 21); *see supra* n. 13.

<sup>29</sup> **Plaintiff’s FN#29:** *Id.* at \*213.

AWS has had knowledge of the '170 patent since at least the time it was served with this Complaint. AWS nevertheless continues to act in wanton disregard of Kove's patent rights.

**RESPONSE:** Denied.

50. Kove is informed and believes, and thereon alleges, that AWS has actively induced the infringement of the '170 patent under 35 U.S.C. § 271(b) by actively inducing the infringing use of the AWS Accused Products by third party users in the United States. Kove is informed and believes, and thereon alleges, that AWS knew or should have known that its conduct would induce others to use the Accused Products in a manner that infringes the '170 patent. Kove is informed and believes, and thereon alleges, that these third parties infringe the '170 patent in violation of 35 U.S.C. § 271(a) by using the AWS Accused Products.

**RESPONSE:** Denied.

51. For example, AWS provides several support websites instructing third parties on the use of the AWS Accused Products, including, without limitation: "Getting Started Resource Center"<sup>30</sup>; "Cloud Solutions"<sup>31</sup>; "AWS Documents"<sup>32</sup>; "AWS Support"<sup>33</sup>; and "AWS Customer Success."<sup>34</sup> These exemplary instructional documentations explain how to use the AWS Accused Products to store and retrieve data. In addition, AWS provides a "step-by-step tutorial [that] will help you store your files in the cloud using Amazon Simple Storage Solution (Amazon S3). Amazon S3 is a service that enables you to store your data (referred to as objects) in at massive scale. In this tutorial, you will create an Amazon S3 bucket, upload a file, retrieve the file and delete the file."<sup>35</sup>

---

<sup>30</sup> **Plaintiff's FN#30:** Getting Started Resource Center, available at <https://aws.amazon.com/getting-started/>.

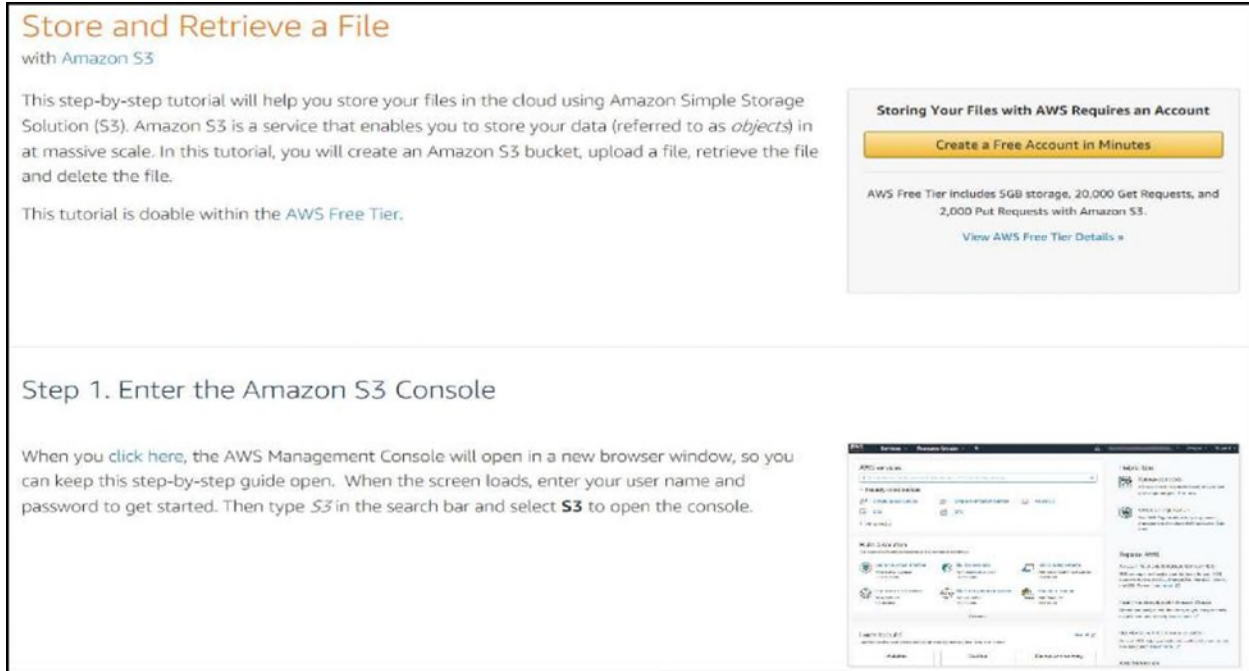
<sup>31</sup> **Plaintiff's FN#31:** Cloud Solutions Find the right solution by application or industry, available at <https://aws.amazon.com/solutions/>.

<sup>32</sup> **Plaintiff's FN#32:** AWS Documentation Find user guides, developer guides, API references, tutorials, and more, available at <https://aws.amazon.com/documentation/>.

<sup>33</sup> **Plaintiff's FN#33:** AWS Support, available at <https://aws.amazon.com/premiumsupport/>.

<sup>34</sup> **Plaintiff's FN#34:** AWS Customer Success, available at <https://aws.amazon.com/solutions/case-studies/>.

<sup>35</sup> **Plaintiff's FN#35:** Store and Retrieve a File with Amazon S3, available at <https://aws.amazon.com/getting-started/tutorials/backup-files-to-amazon-s3/>.



**RESPONSE:** To the extent Paragraph 51 purports to quote documents cited in the footnotes to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Otherwise denied.

52. AWS also provides answers on its website to popular topics, posts, and questions that developers, users, or operators may have about the AWS Accused Products.<sup>36</sup> For instance, AWS News Blog provides updates on Amazon S3.<sup>37</sup> Likewise, AWS educates users about features and developments of the AWS Accused Products: “AWS holds events both online and in-person to bring the cloud computing community together to connect, collaborate, and learn from AWS experts.”<sup>38</sup> Users can also get help by directly asking AWS support developers.<sup>39</sup> AWS further provides instructional videos to teach users how to use the AWS Accused Products.<sup>40</sup>

<sup>36</sup> **Plaintiff’s FN#36:** AWS News Blog, available at <https://aws.amazon.com/blogs>.

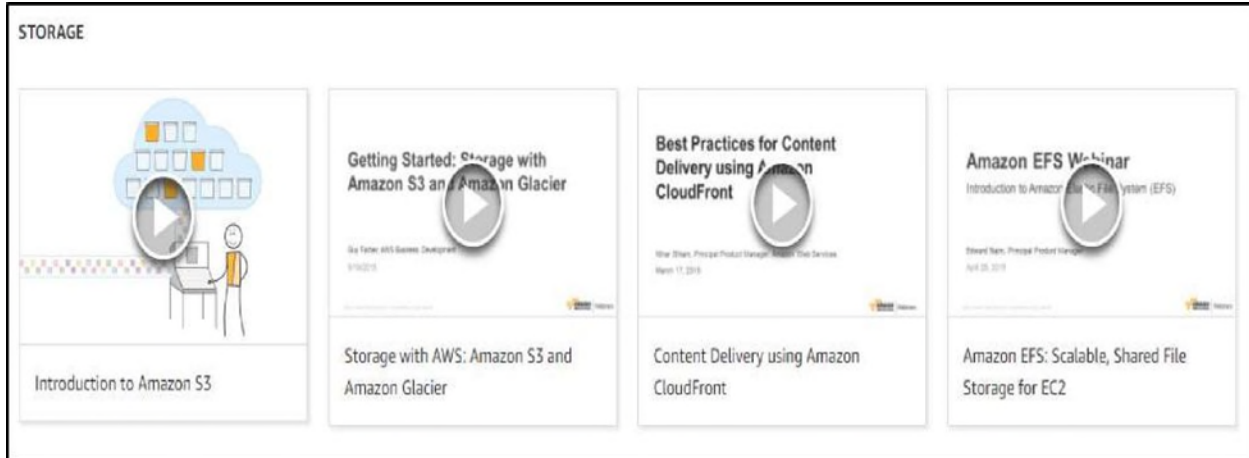
<sup>37</sup> **Plaintiff’s FN#37:** Barr, J., Amazon S3 Update: New Storage Class and General Availability of S3 Select, available at <https://aws.amazon.com/blogs/aws/amazon-s3-update-new-storage-class-general-availability-of-s3-select/> (Attached as Ex. 22).

<sup>38</sup> **Plaintiff’s FN#38:** AWS Events & Webinars available at <https://aws.amazon.com/about-aws/events/>.

<sup>39</sup> **Plaintiff’s FN#39:** See AWS Support, *supra* n. 33.

<sup>40</sup> **Plaintiff’s FN#40:** Getting Started Videos, <https://aws.amazon.com/getting-started/videos/>; Amazon S3, <http://amzn.to/2iNk9IA>; Introduction to Amazon S3, <https://youtu.be/rKpKHulqYOO>.





**RESPONSE:** To the extent Paragraph 52 purports to quote documents and reproduce figures cited in the footnotes to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts or accurately reproduces the cited figure, and refers the Court to those documents for their full content. Otherwise denied.

53. In addition, AWS provides S3 and DynamoDB manuals to help the customer understand, setup, and use the features of the AWS Accused Products, along with user manuals for all AWS products and services, which may be accessed through links at <https://aws.amazon.com/documentation/>. AWS also provides Tools for using the AWS Accused Products, such as SDKs (software development kits), IDE (integrated development environment) Toolkits, and Command Line Tools.<sup>41</sup>

**RESPONSE:** To the extent Paragraph 53 refers to excerpts of documents cited in the footnote to this Paragraph, AWS refers the Court to those documents for their full content. Otherwise denied.

54. As yet another example, in this District, AWS promotes collaborating with “partners and customers” and encourages architects to work “with AWS field sales, pre-sales, training and support teams to help partners and customers learn and use AWS services such as Amazon Elastic Compute Cloud (EC2), Amazon Simple Storage Service (S3), Amazon SimpleDB/RDSdatabases, AWS Identity and Access Management (IAM), etc.”<sup>42</sup> Accordingly, AWS actively induces third parties to infringe the ’170 patent.

<sup>41</sup> **Plaintiff’s FN#41:** Tools for Amazon Web Services, available at <https://aws.amazon.com/tools/>.

<sup>42</sup> **Plaintiff’s FN#42:** Amazon Jobs, Cloud Infrastructure Architect – Chicago – Job ID: 583645 (Attached as Ex. 23).



**RESPONSE:** To the extent Paragraph 54 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

55. Upon information and belief, AWS contributorily infringes the '170 patent under 35 U.S.C. § 271(c) by importing, selling and/or offering to sell within the United States the AWS Accused Products (or components thereof) that constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. For example, the AWS Accused Products, including Amazon S3 and DynamoDB, are material, have no substantial non-infringing uses, and are known by AWS to be especially made or especially adapted for use in the manner claimed in the '170 patent. Accordingly, AWS contributorily infringes the '170 patent.

**RESPONSE:** Denied.

## **COUNT II**

### **AWS'S INFRINGEMENT OF THE '978 PATENT**

56. Kove incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

**RESPONSE:** AWS incorporates herein by reference its responses to Paragraphs 1 through 55.

57. On information and belief, AWS has infringed and will continue to infringe the '978 patent. AWS directly infringes the '978 patent under 35 U.S.C. § 271(a) by making, using, selling, offering for sale, and/or importing in this District and into the United States products and/or methods covered by one or more claims of the '978 patent, including, but not limited to, the AWS Accused Products. As an example, the AWS Accused Products infringe at least claim 17 of the '978 patent.

**RESPONSE:** Denied.

58. Claim 17 is directed to a “method of scaling at least one of capacity and transaction rate capability in a location server in a system having a plurality of location servers for storing and retrieving location information.” Claim 17, for example, recites a method for “providing a transfer protocol configured to transport identifier and location information.” The claimed method also recites “storing location information . . . at a first location server.” Claim 17’s method further requires “receiving an identifier and a location relevant to the identifier at the first location server,” “storing the received location in a location store at the first data location server,” and “transferring a portion of the identifiers and associated locations to a second data location server when a performance criterion of the first location server reaches a predetermined performance limit.”

**RESPONSE:** To the extent Paragraph 58 quotes from excerpts of Claim 17 of the '978 patent, AWS admits that this Paragraph accurately quotes the excerpts, but refers the Court to the '978 patent for its full content. Otherwise denied.

59. On information and belief, AWS provides for a method of scaling at least one of capacity and transaction rate capability in a location server in a system having a plurality of location servers for storing and retrieving location information. For example, DynamoDB manages multiple partitions and servers that support a soft throughput performance limit of 3000 read-capacity units and 1000 write-capacity units. By managing those partitions and servers, DynamoDB provides “some flexibility in your per-partition throughput provisioning by providing burst capacity, as follows. Whenever you are not fully using a partition’s throughput, DynamoDB reserves a portion of that unused capacity for later bursts of throughput to handle usage spikes.”<sup>43</sup>

**RESPONSE:** To the extent Paragraph 59 purports to quote documents cited in the footnote to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Otherwise denied.

60. As explained in previous paragraphs, AWS stores location information at location servers. For example, DynamoDB stores Amazon S3 object identifiers.<sup>44</sup> “The heart of the S3 object index is a DynamoDB table with one item per object, which associates various attributes with the object’s S3 key. Each item contains the S3 key, the size of the object, and any additional attributes to use for lookups.”<sup>45</sup>

**RESPONSE:** To the extent Paragraph 60 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

61. DynamoDB tables organize data location information across servers based on a hash function.<sup>46</sup> For instance, AWS has described DynamoDB as a system that “uses its hash function to determine where to store a new item . . . . Note that the items are not stored in sorted order. Each item’s location is determined by the hash value of its partition key.”<sup>47</sup>

---

<sup>43</sup> **Plaintiff’s FN#43:** See *supra* n. 10 at \*517 (DynamoDB Developer Guide, Ex. 11).

<sup>44</sup> **Plaintiff’s FN#44:** *Id.* at \*523.

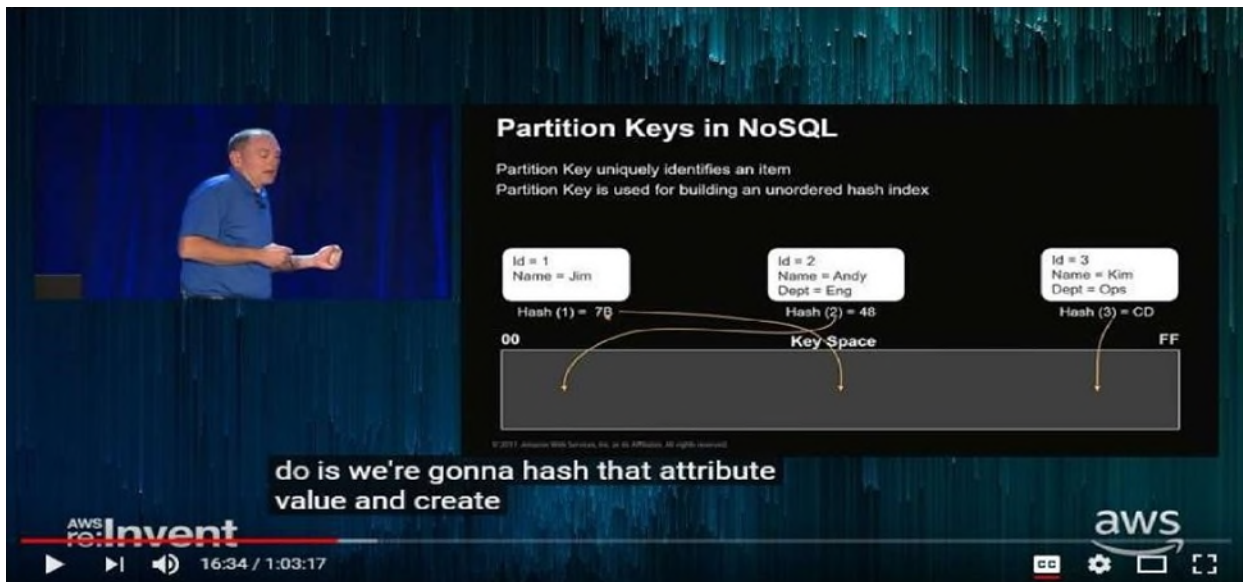
<sup>45</sup> **Plaintiff’s FN#45:** See *supra* n. 21 (Deck, Ex. 19).

<sup>46</sup> **Plaintiff’s FN#46:** See *supra* n. 10 at \*18 (DynamoDB Developer Guide, Ex. 11).

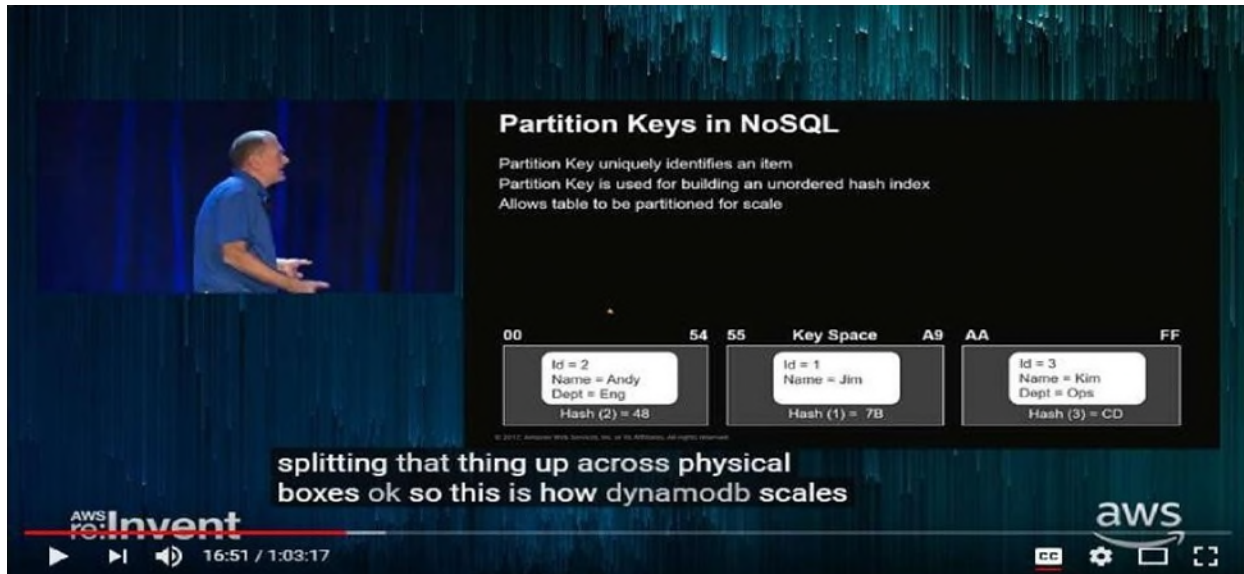
<sup>47</sup> **Plaintiff’s FN#47:** *Id.* at \*18.

**RESPONSE:** To the extent Paragraph 61 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

62. As a further example, it has been explained that DynamoDB hashes location information to create hash tables, which is then split up across a set of servers: “[W]e’re gonna hash that attribute value and create this . . . hash index and lay these items out on an arbitrary key space . . . as you add capacity or you increase the storage and add more items into the table we’re gonna start splitting that thing up across physical boxes ok so this is how DynamoDB scales . . . .”<sup>48</sup>



<sup>48</sup> **Plaintiff’s FN#48:** See *supra* n. 12 (DynamoDB Youtube Video).



**RESPONSE:** To the extent Paragraph 62 quotes from excerpts and reproduces screenshots of the video cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts and reproduces screenshots of the video cited, but refers the Court to that video for its full content. Otherwise denied.

63. In addition, the AWS DynamoDB Developer Guide explains that DynamoDB handles the partition management entirely and the partitions are automatically replicated across multiple Availability Zones within an AWS region.<sup>49</sup> On information and belief, the partition management uses a gossip protocol that allows for “full table routing to other nodes in the system.”<sup>50</sup>

**RESPONSE:** To the extent Paragraph 63 purports to quote documents cited in the footnotes to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content. Otherwise denied.

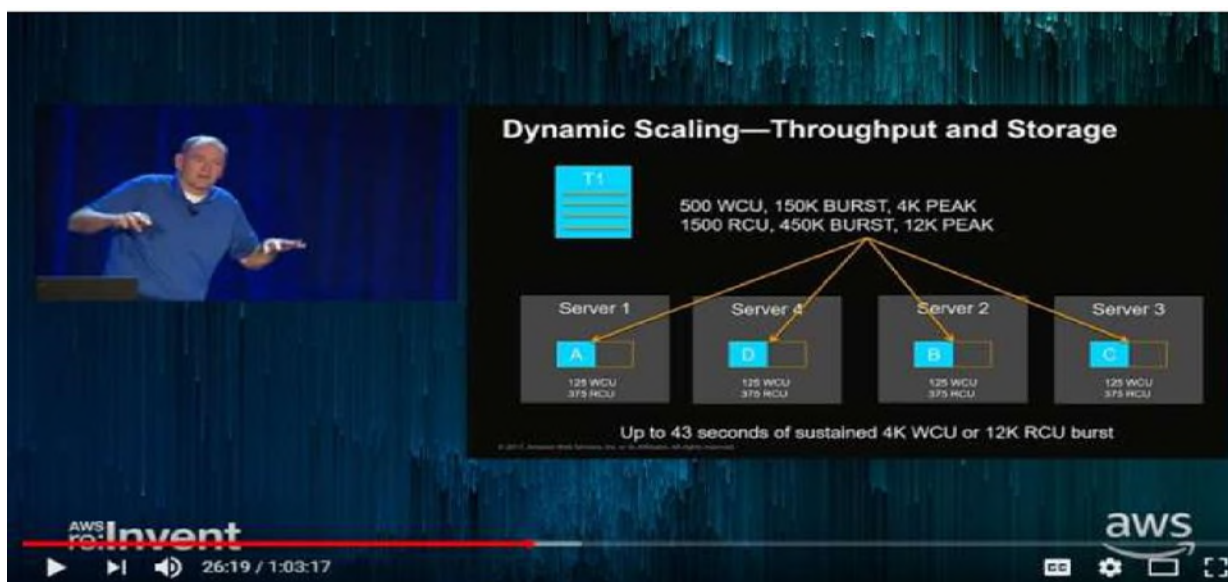
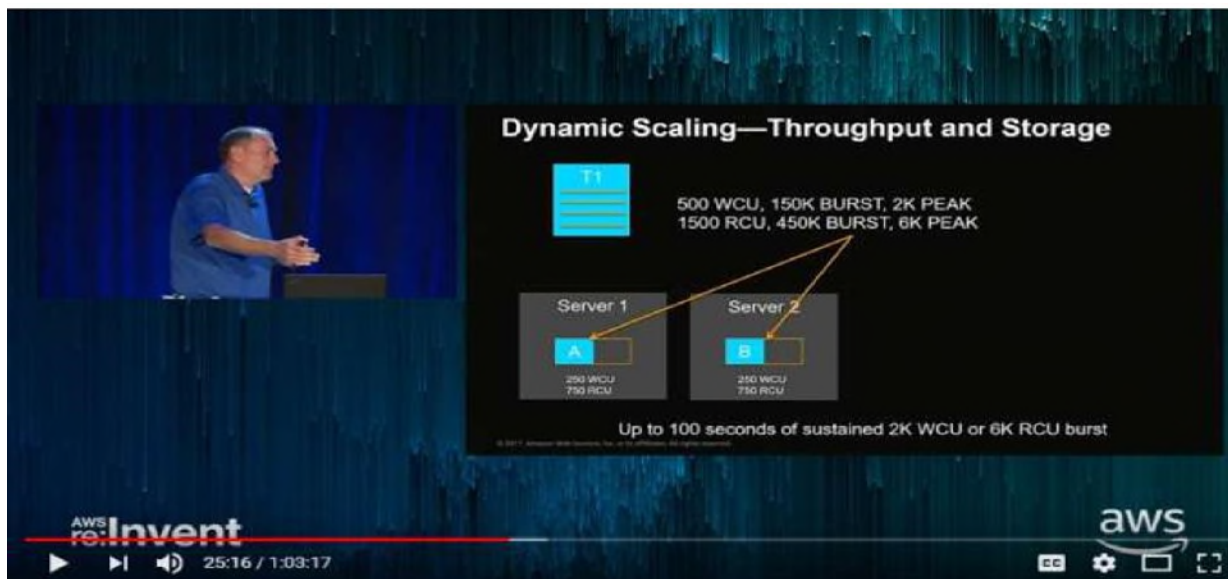
64. Likewise, AWS transfers a portion of the identifiers and associated locations to data location servers when a performance criterion of a location server reaches a predetermined performance limit. DynamoDB’s and Amazon S3’s best practices for optimizing performance depend on the usage request rates—the throughput performance limit. For example, if the workload request rate grows steadily, Amazon S3 automatically partitions the buckets as needed to support higher request rates.<sup>51</sup> As another example, on information and belief, and as an AWS

<sup>49</sup> **Plaintiff’s FN#49:** See *supra* n. 10 at \*18 (DynamoDB Developer Guide, Ex. 11).

<sup>50</sup> **Plaintiff’s FN#50:** See *supra* n. 28 at \*218 (DeCandia, Ex. 21); see also *supra* n. 13.

<sup>51</sup> **Plaintiff’s FN#51:** See *supra* n. 18 at \*535 (S3 Developer Guide, Ex. 17).

senior practice manager explained, DynamoDB may use dynamic scaling or burst capacity—throughput and storage—to automatically spread the data location information across more partitions if the throughput performance limit is exceeded<sup>52</sup>:




<sup>52</sup> **Plaintiff's FN#52:** See *supra* n. 12 (DynamoDB Youtube Video).



**RESPONSE:** To the extent Paragraph 64 reproduces screenshots of the video cited in one of the footnotes to this Paragraph, AWS admits that this Paragraph accurately reproduces those screenshots, but refers the Court to the video for its full content. Otherwise denied.

65. Further still, the AWS DynamoDB Developer Guide shows that the partition data is distributed across multiple servers:



**RESPONSE:** Denied.

66. As depicted, “[t]he large squares represent partitions, and the small squares represent data items in the table. Note that DynamoDB performs partition splits automatically, in the background.”<sup>53</sup>

**RESPONSE:** To the extent Paragraph 66 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Denied that the webpages at the cited links in the footnote to Paragraph 28 contain the quoted language. Otherwise denied.

67. Based on the above, the AWS Accused Products directly infringe at least, but not limited to, claim 17 of the '978 patent.

**RESPONSE:** Denied.

68. AWS also indirectly infringes the '978 patent by inducing others to infringe and/or contributing to the infringement of others, including third party users of the AWS Accused Products in this District and elsewhere in the United States. Specifically, on information and belief,

<sup>53</sup> **Plaintiff's FN#53:** See *supra* n. 10 at \*514 (DynamoDB Developer Guide, Ex. 11).

AWS has had knowledge of the '978 patent since at least the time it was served with this Complaint. AWS nevertheless continues to act in wanton disregard of Kove's patent rights.

**RESPONSE:** Denied.

69. Kove is informed and believes, and thereon alleges, that AWS has actively induced the infringement of the '978 patent under 35 U.S.C. § 271(b) by actively inducing the infringing use of the AWS Accused Products by third party users in the United States. Kove is informed and believes, and thereon alleges, that AWS knew or should have known that its conduct would induce others to use the Accused Products in a manner that infringes the '978 patent. Kove is informed and believes, and thereon alleges, that these third parties infringe the '978 patent in violation of 35 U.S.C. § 271(a) by using the AWS Accused Products. In particular, Kove incorporates by reference paragraphs 51 to 54 as if fully set forth herein.

**RESPONSE:** Denied.

70. Upon information and belief, AWS contributorily infringes the '978 patent under 35 U.S.C. § 271(c) by importing, selling and/or offering to sell within the United States the AWS Accused Products (or components thereof) that constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. For example, the AWS Accused Products, including Amazon S3 and DynamoDB, are material, have no substantial non-infringing uses, and are known by AWS to be especially made or especially adapted for use in manner claimed in the '978 patent. Accordingly, AWS contributorily infringes the '978 patent.

**RESPONSE:** Denied.

### **COUNT III**

#### **AWS'S INFRINGEMENT OF THE '640 PATENT**

71. Kove incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

**RESPONSE:** AWS incorporates herein by reference its responses to Paragraphs 1 through

70.

72. On information and belief, AWS has infringed and will continue to infringe the '640 patent. AWS directly infringes the '640 patent under 35 U.S.C. § 271(a) by making, using, selling, offering for sale, and/or importing in and into the United States products and/or methods covered by one or more claims of the '640 patent, including, but not limited to, the AWS Accused Products. As an example, the AWS Accused Products infringe at least claim 18 of the '640 patent.



**RESPONSE:** Denied.

73. Claim 18 is directed to a “system for retrieving data location information for data stored in a distributed network.” Claim 18 recites, for example, a “data repository configured to store data, wherein the data is associated with an identifier string.” The system also includes a “client responsive to a data query to query a data location server for location information associated with the identifier string.” The “data location server network” according to claim 18 contains “location information associated with the identifier string.” If the “data location server” does not contain the “location information,” “the location server transmits a redirect message to the client, wherein the redirect message contains redirect information for use by the client to calculate a location of a different data location server” containing the location information.

**RESPONSE:** To the extent Paragraph 73 quotes from excerpts of Claim 18 of the ’640 patent, AWS denies that this Paragraph accurately quotes the ’640 patent and refers the Court to the ’640 patent for its full content. Otherwise denied.

74. AWS provides for a system for retrieving data location information for data stored in a distributed network. For example, as shown above in the Architecture Diagram, “[u]sers upload files into Amazon Simple Storage Service (Amazon S3), a highly durable storage infrastructure designed for mission-critical and primary data storage. Amazon S3 makes it easy to store and retrieve any amount of data, at any time. . . . File metadata, version information, and unique identifiers are stored by the application servers on an Amazon DynamoDB table. As the number of files to maintain in the application grows, Amazon DynamoDB tables can store and retrieve any amount of data, and serve any amount of traffic.”<sup>54</sup>

**RESPONSE:** To the extent Paragraph 74 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

75. AWS also provides for a data repository configured to store data, wherein the data is associated with an identifier string. For example, as discussed, Amazon S3 stores data in a files repository. Each data entity stored by Amazon S3, called an Amazon S3 Object, is “uniquely identified within a bucket by a key (name) and a version ID. . . . Amazon S3 can be thought of as a basic data map between ‘bucket + key + version’ and the object itself. Every object in Amazon S3 can be uniquely addressed through the combination of the web service endpoint, bucket name, key, and optionally, a version.”<sup>55</sup>

---

<sup>54</sup> **Plaintiff’s FN#54:** See *supra* n. 17 (Architecture Diagram, Ex. 16).

<sup>55</sup> **Plaintiff’s FN#55:** See *supra* n. 18 at \*3–4 (S3 Developer Guide, Ex. 17).

**RESPONSE:** To the extent Paragraph 75 purports to quote documents cited in the footnote to this Paragraph, AWS denies that the Paragraph accurately quotes the cited excerpts and refers the Court to those documents for their full content.

76. Further, as discussed, Amazon S3 stores location information for S3 objects in a DynamoDB table. On information and belief, DynamoDB tables are distributed across multiple servers. AWS has explained, for example, that “DynamoDB stores data in partitions. A partition is an allocation of storage for a table, backed by solid-state drives (SSDs) and automatically replicated across multiple Availability Zones within an AWS Region. Partition management is handled entirely by DynamoDB—you never have to manage partitions yourself.”<sup>56</sup>

**RESPONSE:** To the extent Paragraph 76 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

77. Likewise, location information for Amazon S3 objects is stored in DynamoDB, along with the objects’ metadata. The Amazon S3 “index subsystem manages the metadata and location information of all S3 objects in the region,”<sup>57</sup> and it has been explained that “[t]he heart of the S3 object index is a DynamoDB table with one item per object, which associates various attributes with the object’s S3 key. Each item contains the S3 key, the size of the object, and any additional attributes to use for lookups.”<sup>58</sup> AWS’s “Frequently Asked Questions About Amazon DynamoDB” explains that “file pointers (possibly to Amazon S3 objects) are best saved in Amazon DynamoDB,” as opposed to in Amazon S3.<sup>59</sup>

**RESPONSE:** To the extent Paragraph 77 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

78. On information and belief, DynamoDB tables organize data location information across servers based on a hash function.<sup>60</sup> For instance, AWS has described DynamoDB as a system that “uses its hash function to determine where to store a new item . . . . Note that the items

---

<sup>56</sup> **Plaintiff’s FN#56:** See *supra* n. 10 at \*18 (DynamoDB Developer Guide, Ex. 11).

<sup>57</sup> **Plaintiff’s FN#57:** See *supra* n. 19 (Summary of the Amazon S3 Service Disruption in the Northern Virginia, Ex. 18).

<sup>58</sup> **Plaintiff’s FN#58:** See *supra* n. 21 (Deck, Ex. 19).

<sup>59</sup> **Plaintiff’s FN#59:** See *supra* n. 22 at \*4–5 (Amazon DynamoDB FAQs, Ex. 20).

<sup>60</sup> **Plaintiff’s FN#60:** See *supra* n. 10 at \*18 (DynamoDB Developer Guide, Ex. 11).

are not stored in sorted order. Each item's location is determined by the hash value of its partition key."<sup>61</sup>

**RESPONSE:** To the extent Paragraph 78 quotes from excerpts of documents cited in the footnotes to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

79. Further, Amazon S3 includes the AWS Management Console,<sup>62</sup> which responds to requests to download Amazon S3 objects by querying the locations of the requested objects.<sup>63</sup>

**RESPONSE:** To the extent Paragraph 79 refers to excerpts of documents cited in the footnotes to this Paragraph, AWS refers the Court to those documents for their full content. Otherwise denied.

80. Further still, AWS provides for a redirect message that contains redirect information for use to calculate a location of a different data location server. For example, Amazon S3 supports request redirection, as illustrated below, such that if "a request arrives at the wrong Amazon S3 location, Amazon S3 responds with a temporary redirect that tells the requester to send the request to a new endpoint."<sup>64</sup>

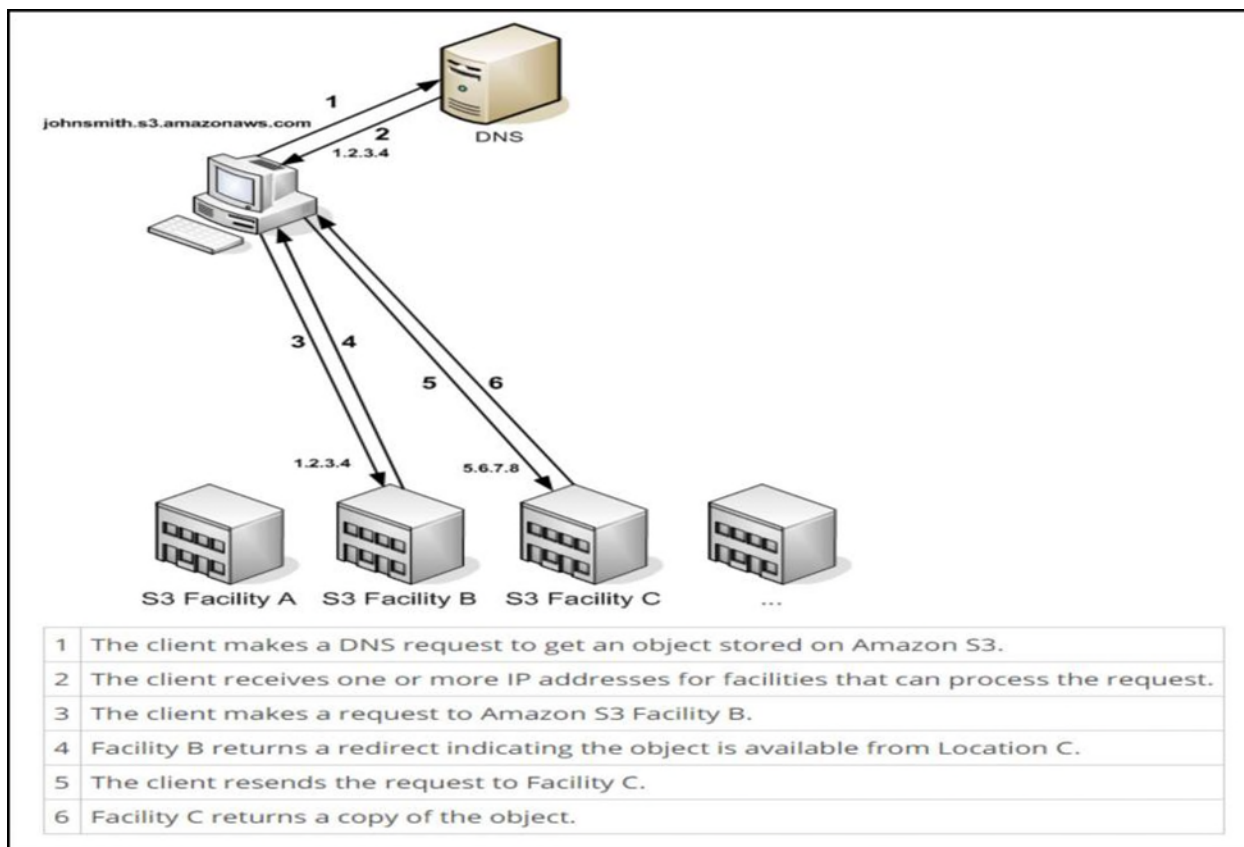
---

<sup>61</sup> **Plaintiff's FN#61:** *Id.*

<sup>62</sup> **Plaintiff's FN#62:** AWS Management Console, Getting Started with the AWS Management Console, at \*1, \*9, available at <https://docs.aws.amazon.com/awsconsolehelpdocs/latest/gsg/console-help-gsg.pdf> (Attached as Ex. 24).

<sup>63</sup> **Plaintiff's FN#63:** Amazon Simple Storage Service Console User Guide, Uploading, Downloading, and Managing Objects, at \*31–41, available at <https://docs.aws.amazon.com/AmazonS3/latest/user-guide/s3-user-guide.pdf> (While the full document is incorporated by reference herein, an excerpted version of the document is attached as Ex. 25); *see also* Amazon's Web Version of the S3 Console User Guide at <https://docs.aws.amazon.com/AmazonS3/latest/user-guide/what-is-s3.html>.

<sup>64</sup> **Plaintiff's FN#64:** *See supra* n. 18 at \*530–531 (S3 Developer Guide, Ex. 17).



**RESPONSE:** To the extent Paragraph 80 quotes from excerpts and reproduces figures of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts and reproduces the cited figure, but refers the Court to those documents for their full content. Otherwise denied.

81. Moreover, on information and belief, AWS receives get and put operations as discussed above: “Both get and put operations are invoked using Amazon’s infrastructure-specific request processing framework over HTTP. There are two strategies that a client can use to select a node: (1) route its request through a generic load balancer that will select a node based on load information, or (2) use a partition-aware client library that routes requests directly to the appropriate coordinator nodes. . . . A node handling a read or write operation is known as the coordinator. Typically, this is the first among the top N nodes in the preference list. If the requests are received through a load balancer, requests to access a key may be routed to any random node in the ring. In this scenario, the node that receives the request will not coordinate it if the node is not in the top N of the requested key’s preference list. Instead, that node will forward the request to the first among the top N nodes in the preference list. . . . Read and write operations involve the first N healthy nodes in the preference list, skipping over those that are down or inaccessible. When all nodes are healthy, the top N nodes in a key’s preference list are accessed. When there

are node failures or network partitions, nodes that are lower ranked in the preference list are accessed.”<sup>65</sup>

**RESPONSE:** To the extent Paragraph 81 quotes from excerpts of documents cited in the footnote to this Paragraph, AWS admits that this Paragraph accurately quotes the cited excerpts, but refers the Court to those documents for their full content. Otherwise denied.

82. Based on the above, the AWS Accused Products directly infringe at least, but not limited to, claim 18 of the '640 patent.

**RESPONSE:** Denied.

83. AWS also indirectly infringes the '640 patent by inducing others to infringe and/or contributing to the infringement of others, including third party users of the AWS Accused Products in this District and elsewhere in the United States. Specifically, on information and belief, AWS has had knowledge of the '640 patent since at least the time it was served with this Complaint. AWS nevertheless continues to induce others to infringe in wanton disregard of Kove's patent rights.

**RESPONSE:** Denied.

84. Kove is informed and believes, and thereon alleges, that AWS has actively induced the infringement of the '640 patent under 35 U.S.C. § 271(b) by actively inducing the infringing use of the AWS Accused Products by third party users in the United States. Kove is informed and believes, and thereon alleges, that AWS knew or should have known that its conduct would induce others to use the Accused Products in a manner that infringes the '640 patent. Kove is informed and believes, and thereon alleges, that these third parties infringe the '640 patent in violation of 35 U.S.C. § 271(a) by using the AWS Accused Products. In particular, Kove incorporates by reference paragraphs 51 to 54 as if fully set forth herein.

**RESPONSE:** Denied.

85. Upon information and belief, AWS contributorily infringes the '640 patent under 35 U.S.C. § 271(c) by importing, selling and/or offering to sell within the United States the AWS Accused Products (or components thereof) that constitute a material part of the claimed invention and are not staple articles of commerce suitable for substantial non-infringing use. For example, the AWS Accused Products, including Amazon S3 and DynamoDB, are material, have no substantial non-infringing uses, and are known by AWS to be especially made or especially adapted for use in manner claimed in the '640 patent. Accordingly, AWS contributorily infringes the '640 patent.

---

<sup>65</sup> **Plaintiff's FN#65:** See *supra* n. 28 at \*211 (DeCandia, Ex. 21); see also *supra* n. 13.

**RESPONSE:** Denied.

**DAMAGES**

86. Kove incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

**RESPONSE:** AWS incorporates herein by reference its responses to Paragraphs 1 through 85.

87. As a result of AWS's acts of infringement, Kove has suffered actual and consequential damages; however, Kove does not yet know the full extent of the infringement and its extent cannot be ascertained except through discovery and special accounting. To the fullest extent permitted by law, Kove seeks recovery of damages at least for reasonable royalties, unjust enrichment, lost profits, and/or benefits received by AWS as a result of infringing the patents-in-suit. Kove further seeks any other damages to which Kove is entitled under law or in equity.

**RESPONSE:** Denied.

**IRREPARABLE HARM TO KOVE**

88. Kove incorporates by reference the foregoing paragraphs of this Complaint as if fully set forth herein.

**RESPONSE:** AWS incorporates herein by reference its responses to Paragraphs 1 through 87.

89. Kove has been irreparably harmed by AWS's acts of infringement and will continue to be irreparably harmed unless and until AWS's acts of infringement are enjoined by this Court. Kove has no adequate remedy at law to redress AWS's continuing acts of infringement, and money damages will not suffice to remedy the harms to Kove. The hardships that would be imposed upon AWS by an injunction are less than those faced by Kove should an injunction not issue. Furthermore, the public interest would be served by issuance of an injunction.

**RESPONSE:** Denied.

**ATTORNEYS' FEES**

90. Kove is entitled to recover reasonable and necessary attorneys' fees under applicable law.

**RESPONSE:** Denied.

**DEMAND FOR JURY TRIAL**

91. Kove hereby demands a jury trial on its claims for patent infringement.

**RESPONSE:** This paragraph does not require a response. To the extent a response is required, AWS also demands a jury trial on Plaintiff's claims and AWS's counterclaims.

**PRAYER FOR RELIEF**

WHEREFORE, Kove respectfully requests that this Court enter judgment in its favor and grant the following relief:

A. A judgment that the AWS Accused Products directly and/or indirectly infringe the '170, '978, and '640 patents;

B. That such infringement is willful;

C. An order permanently enjoining AWS and its respective officers, directors, agents, partners, servants, employees, attorneys, licensees, successors, and assigns, and those in active concert or participation with any of them, from engaging in infringing activities with respect to the '170, '978, and '640 patents;

D. A ruling finding that this case is exceptional and awarding Kove its reasonable attorneys' fees under 35 U.S.C. § 285;

E. A judgment and order requiring AWS to pay Kove's damages in an amount adequate to compensate Kove for AWS's infringement, but in no event less than a reasonable royalty under 35 U.S.C. § 284, including supplemental damages for any continuing post-verdict infringement up until entry of judgment and beyond, with accounting, as needed;

F. An award of enhanced damages pursuant to 35 U.S.C. § 284;

G. In the alternative, in the event injunctive relief is not granted as requested by Kove, an award of a mandatory future royalty payable on each future product sold by AWS that is found to infringe one or more claims of the '170, '978, and '640 patents, and on all future products which are not colorably different from products found to infringe;

H. A judgment and order requiring AWS to pay Kove's costs of this action (including all disbursements);

I. An order for accounting of damages;

J. A judgment and order requiring AWS to pay pre-judgment and post-judgment interest to the full extent allowed under the law; and,

K. Such other and further relief as the Court may deem just and proper under the circumstances.

**RESPONSE:** AWS denies all allegations included in Plaintiff's prayer for relief and further denies that Plaintiff is entitled to any relief on any claim.



### **ADDITIONAL DEFENSES**

AWS alleges and asserts the following defenses in response to the allegations, undertaking the burden of proof only as to those defenses deemed affirmative defenses by law, regardless of how such defenses are denominated herein. In addition to the defenses described below, AWS specifically reserves all rights to allege additional defenses that become known through the course of discovery.

#### **FIRST ADDITIONAL DEFENSE: NON-INFRINGEMENT**

AWS does not make, use, sell, offer for sale, or import into the United States, and has not made, used, sold, offered for sale or imported into the United States any products or services that infringe any valid claim of the patents-in-suit willfully, directly, indirectly, contributorily, through the doctrine of equivalents, or otherwise, and has not induced others to infringe any valid claim of the patents-in-suit. AWS incorporates by reference the allegations made in its pleadings and filings in this case, including its Initial and Final Responses to Kove's Infringement Contentions.

#### **SECOND ADDITIONAL DEFENSE: INVALIDITY**

The claims of the patents-in-suit are invalid for failure to meet the conditions for patentability set forth in Title 35 of the United States Code, including, but not limited to, 35 U.S.C. §§ 100 *et seq.*, 101, 102, 103, 112, the requirements of the Code of Federal Regulations, as well as general principles of patent law. AWS incorporates by reference the allegations made in its pleadings and filings in this case, including its Initial and its Final Invalidity Contentions, its Motion to Dismiss and its Counterclaims. *See, e.g.*, Dkts. 38, 45; *infra* at 1–88.

#### **THIRD ADDITIONAL DEFENSE: FAILURE TO STATE A CLAIM**

Plaintiff fails to state a claim upon which relief can be granted.

**FOURTH ADDITIONAL DEFENSE: UNENFORCEABILITY**

Plaintiff's attempted enforcement of the patents-in-suit against AWS is barred by the doctrine of inequitable conduct.

On information and belief, Mr. Kent Genin of Brinks Hofer Gilson & Lione (the "prosecuting attorney") filed and prosecuted the patent applications that ultimately issued as the '978, '640, and '170 patents. On information and belief, Mr. Genin is a patent attorney registered with the United States Patent and Trademark Office ("USPTO"). Mr. John K. Overton and Mr. Stephen W. Bailey (the "Applicants") are listed as the named inventors of the patents-in-suit. On information and belief, the Applicants were involved in and knowledgeable about the prosecution of the patents-in-suit.

The prosecuting attorney and the Applicants have "a duty of candor and good faith in dealing with the Office [USPTO], which includes a duty to disclose to the Office all information known to that individual to be material to patentability." 37 C.F.R. § 1.56. On information and belief, Applicants were aware of the paper "Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web," by David Karger et al. (the "Karger paper") during prosecution of the applications that became the patents-in-suit. Kove produced the Karger paper during discovery in this litigation, and, on information and belief, the Karger paper was in files associated with at least one of the Applicants since January 3, 1999 and during the prosecution of the applications that became the patents-in-suit. Kove identified Mr. Overton as the custodian for this document.

The Karger paper is highly material prior art to the patents-in-suit, but neither the prosecuting attorney nor the Applicants disclosed the Karger paper to the USPTO. The Karger paper is material to the patents-in-suit at least because it describes "caching protocols for

distributed networks that can be used to decrease or eliminate the occurrence of hot spots” in networks such as the World Wide Web. The Karger paper is cited in AWS’s Final Invalidity Contentions, which are hereby incorporated by reference, as prior art that “either anticipate[s] or render[s] obvious the patents-in-suit in combination with one or more other references and/or the knowledge of one of ordinary skill in the art around the time of the invention.”

On information and belief, if the Karger paper were before the USPTO during prosecution of the applications that became the patents-in-suit, at least one claim of each of the patents-in-suit would not have issued. More specifically, at least claims 1, 3, 6, 10, 14, 17, 23, 24 and 31 of the ’978 patent would not have issued had the Karger paper been before the USPTO during prosecution; at least claims 1, 2, 6, 8, 12 and 15 of the ’170 patent would not have issued had the Karger paper been before the USPTO during prosecution; and at least claims 17, 18 and 24 of the ’640 patent would not have issued had the Karger paper been before the USPTO during prosecution.

On information and belief, Applicants and/or the prosecuting attorney made a deliberate decision to withhold the Karger paper from the USPTO with knowledge of the Karger paper’s materiality and with the specific intent to deceive the USPTO. On information and belief, the patents-in-suit are unenforceable because Applicants and/or the prosecuting attorney failed to disclose the Karger paper and made material misrepresentations about the prior art to the USPTO.

**FIFTH ADDITIONAL DEFENSE: FAILURE TO MARK**

Plaintiff and any predecessors in interest to the patents-in-suit failed to properly mark any of their relevant products or materials as required by 35 U.S.C. § 287, or otherwise to give proper notice that AWS’s actions allegedly infringe the patents-in-suit, and AWS is not liable to

Plaintiff for the acts alleged to have been performed before it received actual notice that it was allegedly infringing the patents-in-suit.

**SIXTH ADDITIONAL DEFENSE: BAR TO DAMAGES**

Plaintiff's claims for damages are barred, in whole or in part, by the six-year limitation of damages in 35 U.S.C. § 286.

**SEVENTH ADDITIONAL DEFENSE: EQUITABLE DEFENSES**

Plaintiff's attempted enforcement of the patents-in-suit against AWS is barred by one or more of the equitable doctrines of laches, estoppel, acquiescence, waiver, unclean hands, and/or other equitable principles.

Amazon Simple Storage Service ("S3") was released in March 2006, and DynamoDB was released in January 2012. But Plaintiff did not file the instant suit until December 2018, nearly 7 years after the release of DynamoDB, and over 12.5 years after the release S3. Further, Plaintiff's Complaint extensively discusses Amazon documents that put Plaintiff on notice of Amazon's alleged infringement years before Plaintiff filed this litigation. As but a few examples, Plaintiff's Complaint cites to Amazon's DynamoDB Developer Guide 13 times; that guide was released in August 2012. Plaintiff's Complaint cites Amazon's S3 Developer Guide 4 times; that guide was released in March 2006. Plaintiff additionally cites articles about Amazon's services from 2007, 2012, 2013, and 2014.

Further, Plaintiff failed to mark any of its products in accordance with 35 U.S.C. § 287, or to otherwise give any notice that AWS's actions allegedly infringe the patents-in-suit, meaning that AWS became aware of Plaintiff's allegations only when the instant litigation was filed in December 2018. Across the 12.5 years of operation of its S3 service, and the nearly 7 years of its DynamoDB service, AWS reasonably relied on Plaintiff's delay and inaction when it

continued to supply its S3 and DynamoDB products and otherwise grow its cloud storage services. AWS is significantly prejudiced by Plaintiff's delay, as years earlier, among other options, AWS could have chosen to cease providing its services, to obtain licenses for the patents-in-suit, or to file a declaratory judgment action.

To the extent Plaintiff seeks equitable relief, the doctrine of laches bars such relief because of Plaintiff's delay, coupled with Amazon's reasonable reliance and prejudice. Further, Plaintiff's conduct demonstrates an intentional relinquishment of a known right, and active consent to AWS's conduct, constituting waiver and acquiescence.

Plaintiff's claims are also barred by the equitable doctrine of unclean hands because Plaintiff acted in bad faith with respect to the subject of this lawsuit, including in its conduct of discovery in this case. AWS served its First Set of Requests for Production of documents on June 20, 2019. Among other things, those requests sought documents related to the prosecution of the Asserted Patents and Plaintiff's knowledge of prior art. Plaintiff produced few documents responsive to those requests. Its document productions consisted largely of publicly-available correspondence with the Patent Office and publicly-available open source code. Throughout the fall of 2019, AWS sent Plaintiff multiple discovery letters detailing the deficiencies in Plaintiff's productions and the parties conferred several times concerning AWS's requests that Plaintiff produce all responsive documents.

Finally, on February 10—only three days before AWS's deadline to serve its Final Invalidity Contentions—Plaintiff produced over 20,000 pages of documents to AWS. Plaintiff made two additional productions totaling over 19,000 pages in the days immediately following that deadline. The February productions contained many documents relating to the prosecution of the patents-in-suit. Those documents are highly relevant to several issues in this case,

including, among other things, a defense of inequitable conduct. Although AWS requested the basis for this delay, Plaintiff did not offer any explanation for waiting over seven months after AWS served its document requests to produce those documents.

Because of Plaintiff's discovery misconduct—coupled with the patent applicants' inequitable conduct in prosecution of the patents-in-suit, as set forth in AWS's fourth additional defense on unenforceability—the doctrine of unclean hands bars Plaintiff's claims.

**EIGHTH ADDITIONAL DEFENSE: DOUBLE-PATENTING**

Claims 1, 3, 6, 10, 14, 17, 23, 24, 30 and 31 of the '978 patent are invalid under the doctrine of obviousness-type double patenting.

**NINTH ADDITIONAL DEFENSE: LACK OF STANDING**

Plaintiff lacks standing to bring suit under the counts alleged in the Complaint.



## **COUNTERCLAIMS**

Counterclaimant Amazon Web Services, Inc. (“AWS”), for its Counterclaims against Kove IO, Inc., states as follows:

### **THE PARTIES**

1. AWS is a corporation existing under the laws of Delaware with a principal place of business at 410 Terry Avenue, North Seattle, Washington 98109.

2. On information and belief, Kove IO, Inc. (“Kove”) is a corporation existing under the laws of Delaware with a principal place of business at 14 North Peoria Street, Suit 2H, Chicago, Illinois 60607.

### **JURISDICTION AND VENUE**

3. This is an action under the patent laws of the United States. This Court has subject matter jurisdiction over these counterclaims under at least 28 U.S.C. §§ 1331 & 1338(a).

4. Kove (“Counter-Defendant”) accuses AWS of infringing three patent claims in its Complaint, and in its Infringement Contentions has accused AWS of infringing additional claims: claims 1, 3, 6, 10, 14, 17, 23, 24, 30 and 31 of U.S. Patent No. 7,233,978; claims 1, 2, 6, 8, 12 and 15 of U.S. Patent No. 7,814,170; and claims 17, 18 and 24 of U.S. Patent No. 7,103,640. By filing its Complaint, Counter-Defendant has submitted to personal jurisdiction in this District.

5. Kove has consented to venue under 28 U.S.C. §§ 1391 and 1400 by filing the original action.

### **FIRST COUNTERCLAIM: NON-INFRINGEMENT OF U.S. PATENT NO. 7,233,978**

6. The allegations of paragraphs 1 through 5 are realleged and reincorporated by reference as if fully set forth herein.

7. Counter-Defendant claims to be the owner of U.S. Patent No. 7,233,978 (“the ’978 patent”).

8. Counter-Defendant contends, or has contended, that AWS infringes claims of the ’978 patent.

9. AWS does not infringe and has not infringed the ’978 patent directly, indirectly, contributorily, by inducement, by equivalence, or in any other way.

10. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the non-infringement of the ’978 patent.

11. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the ’978 patent are not infringed.

**SECOND COUNTERCLAIM: INVALIDITY OF U.S. PATENT NO. 7,233,978**

12. The allegations of paragraphs 1 through 11 are realleged and reincorporated by reference as if fully set forth herein.

13. The ’978 patent is invalid for failure to meet the conditions for patentability set forth in Title 35 of the United States Code, including, but not limited to, 35 U.S.C. § 101 *et seq.*, the requirements of the Code of Federal Regulations, as well as general principles of patent law.

14. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the invalidity of the ’978 patent.

15. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the ’978 patent are invalid.

**THIRD COUNTERCLAIM: UNENFORCEABILITY OF U.S. PATENT NO. 7,233,978**

16. The allegations of paragraphs 1 through 15 are realleged and reincorporated by reference as if fully set forth herein.

17. On information and belief, Mr. Kent Genin of Brinks Hofer Gilson & Lione (the “prosecuting attorney”) filed and prosecuted U.S. Patent Application No. 09/872,736 (“the ’736 application”), the patent application that ultimately issued as the ’978 patent.

18. On information and belief, Mr. Genin is a patent attorney registered with the United States Patent and Trademark Office (“USPTO”).

19. Mr. John K. Overton and Mr. Stephen W. Bailey (the “Applicants”) are listed as the named inventors of the ’978 patent.

20. On information and belief, the Applicants were involved in and knowledgeable about the prosecution of the ’978 patent.

21. The prosecuting attorney and the Applicants have “a duty of candor and good faith in dealing with the Office [USPTO], which includes a duty to disclose to the Office all information known to that individual to be material to patentability.” 37 C.F.R. § 1.56.

22. On information and belief, Applicants and/or the prosecuting attorney were aware of the paper “Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web,” by David Karger et al. (the “Karger paper”) during prosecution of the ’736 application. Kove produced the Karger paper during discovery in this litigation, and, on information and belief, the Karger paper was in the same file since January 3, 1999 and during the prosecution of the ’736 application. Kove identified Mr. Overton as the custodian for this document.

23. The Karger paper is highly material prior art to the ’978 patent, but neither the prosecuting attorney nor the Applicants disclosed the Karger paper to the USPTO.

24. The Karger paper is material to the ’978 patent at least because it describes “caching protocols for distributed networks that can be used to decrease or eliminate the occurrence of hot

spots” in networks such as the World Wide Web. The Karger paper is cited in AWS’s Final Invalidity Contentions, which are hereby incorporated by reference, as prior art that “either anticipate[s] or render[s] obvious the patents-in-suit in combination with one or more other references and/or the knowledge of one of ordinary skill in the art around the time of the invention.”

25. On information and belief, if the Karger paper were before the USPTO during prosecution of the ’736 application, at least one claim of the ’978 patent would not have issued. More specifically, at least claims 1, 3, 6, 10, 14, 17, 23, 24, 30 and 31 would not have been issued had the Karger paper been before the USPTO during prosecution.

26. On information and belief, Applicants and/or the prosecuting attorney made a deliberate decision to withhold the Karger paper from the USPTO with knowledge of the Karger paper’s materiality and with the specific intent to deceive the USPTO.

27. On information and belief, the ’978 patent is unenforceable because Applicants and/or the prosecuting attorney failed to disclose the Karger paper and made material misrepresentations about the prior art to the USPTO.

28. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the ’978 patent are unenforceable for the reasons set forth paragraphs 16–27 above.

**FOURTH COUNTERCLAIM: UNPATENTABILITY OF U.S. PATENT NO. 7,233,978**

29. The allegations of paragraphs 1 through 28 are realleged and reincorporated by reference as if fully set forth herein.

30. The ’978 patent does not claim patentable subject matter under 35 U.S.C. § 101 and is thus unpatentable.

31. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to unpatentability of the '978 patent.

32. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '978 patent are unpatentable under 35 U.S.C. § 101.

**FIFTH COUNTERCLAIM: NON-INFRINGEMENT OF U.S. PATENT NO. 7,814,170**

33. The allegations of paragraphs 1 through 32 are realleged and reincorporated by reference as if fully set forth herein.

34. The Counter-Defendant claims to be the owner of U.S. Patent No. 7,814,170 (“the '170 patent”).

35. Counter-Defendant contends, or has contended, that AWS infringes claims of the '170 patent.

36. AWS does not infringe and has not infringed the '170 patent directly, indirectly, contributorily, by inducement, by equivalence, or in any other way.

37. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the non-infringement of the '170 patent.

38. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '170 patent are not infringed.

**SIXTH COUNTERCLAIM: INVALIDITY OF U.S. PATENT NO. 7,814,170**

39. The allegations of paragraphs 1 through 38 are realleged and reincorporated by reference as if fully set forth herein.

40. The '170 patent is invalid for failure to meet the conditions for patentability set forth in Title 35 of the United States Code, including, but not limited to, 35 U.S.C. § 101 *et seq.*, the requirements of the Code of Federal Regulations, as well as general principles of patent law.

41. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the invalidity of the '170 patent.

42. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '170 patent are invalid.

**SEVENTH COUNTERCLAIM: UNENFORCEABILITY OF U.S. PATENT NO. 7,814,170**

43. The allegations of paragraphs 1 through 42 are realleged and reincorporated by reference as if fully set forth herein.

44. On information and belief, Mr. Kent Genin of Brinks Hofer Gilson & Lione (the “prosecuting attorney”) filed and prosecuted U.S. Patent Application No. 11/354,224 (“the '224 application”), the patent application that ultimately issued as the '170 patent.

45. On information and belief, Mr. Genin is a patent attorney registered with the USPTO.

46. Mr. John K. Overton and Mr. Stephen W. Bailey (the “Applicants”) are listed as the named inventors of the '170 patent.

47. On information and belief, the Applicants were involved in and knowledgeable about the prosecution of the '170 patent.

48. The prosecuting attorney and the Applicants have “a duty of candor and good faith in dealing with the Office [USPTO], which includes a duty to disclose to the Office all information known to that individual to be material to patentability.” 37 C.F.R. § 1.56.



49. On information and belief, Applicants and/or the prosecuting attorney were aware of the paper “Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web,” by David Karger et al. (the “Karger paper”) during prosecution of the ’224 application. Kove produced the Karger paper during discovery in this litigation, and, on information and belief, the Karger paper was in the same file since January 3, 1999 and during the prosecution of the ’224 application. Kove identified Mr. Overton as the custodian for this document.

50. The Karger paper is highly material prior art to the ’170 patent, but neither the prosecuting attorney nor the Applicants disclosed the Karger paper to the USPTO. More specifically, at least claims 1, 2, 6, 8, 12 and 15 would not have issued had the Karger paper been before the USPTO during prosecution.

51. The Karger paper is material to the ’170 patent at least because it describes “caching protocols for distributed networks that can be used to decrease or eliminate the occurrence of hot spots” in networks such as the World Wide Web. The Karger paper is cited in AWS’s Final Invalidity Contentions, which are hereby incorporated by reference, as prior art that “either anticipate[s] or render[s] obvious the patents-in-suit in combination with one or more other references and/or the knowledge of one of ordinary skill in the art around the time of the invention.”

52. On information and belief, Applicants and/or the prosecuting attorney made a deliberate decision to withhold the Karger paper from the USPTO with knowledge of the Karger paper’s materiality and with the specific intent to deceive the USPTO.

53. On information and belief, the ’170 patent is unenforceable because Applicants and/or the prosecuting attorney failed to disclose the Karger paper and made material misrepresentations about the prior art to the USPTO.

54. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '170 patent are unenforceable for the reasons set forth paragraphs 43–53 above.

**EIGHTH COUNTERCLAIM: UNPATENTABILITY OF U.S. PATENT NO. 7,814,170**

55. The allegations of paragraphs 1 through 54 are realleged and reincorporated by reference as if fully set forth herein.

56. The '170 patent does not claim patentable subject matter under 35 U.S.C. § 101 and is thus unpatentable.

57. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to unpatentability of the '170 patent.

58. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '170 patent are unpatentable under 35 U.S.C. § 101.

**NINTH COUNTERCLAIM: NON-INFRINGEMENT OF U.S. PATENT NO. 7,103,640**

59. The allegations of paragraphs 1 through 58 are realleged and reincorporated by reference as if fully set forth herein.

60. The Counter-Defendant claims to be the owner of U.S. Patent No. 7,103,640 (“the '640 patent”).

61. Counter-Defendant contends, or has contended, that AWS infringes claims of the '640 patent.

62. AWS does not infringe and has not infringed the '640 patent directly, indirectly, contributorily, by inducement, by equivalence, or in any other way.

63. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the non-infringement of the '640 patent.

64. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '640 patent are not infringed.

**TENTH COUNTERCLAIM: INVALIDITY OF U.S. PATENT NO. 7,103,640**

65. The allegations of paragraphs 1 through 64 are realleged and reincorporated by reference as if fully set forth herein.

66. The '640 patent is invalid for failure to meet the conditions for patentability set forth in Title 35 of the United States Code, including, but not limited to, 35 U.S.C. § 101 *et seq.*, the requirements of the Code of Federal Regulations, as well as general principles of patent law.

67. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the invalidity of the '640 patent.

68. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '640 patent are invalid.

**ELEVENTH COUNTERCLAIM:  
UNENFORCEABILITY OF U.S. PATENT NO. 7,103,640**

69. The allegations of paragraphs 1 through 68 are realleged and reincorporated by reference as if fully set forth herein.

70. On information and belief, Mr. Kent Genin of Brinks Hofer Gilson & Lione (the "prosecuting attorney") filed and prosecuted U.S. Patent Application No. 09/661,222 ("the '222 application"), the patent application that ultimately issued as the '640 patent.

71. On information and belief, Mr. Genin is a patent attorney registered with the United States Patent and Trademark Office.

72. Mr. John K. Overton and Mr. Stephen W. Bailey (the “Applicants”) are listed as the named inventors of the ’640 patent.

73. On information and belief, the Applicants were involved in and knowledgeable about the prosecution of the ’640 patent.

74. The prosecuting attorney and the Applicants have “a duty of candor and good faith in dealing with the Office [USPTO], which includes a duty to disclose to the Office all information known to that individual to be material to patentability.” 37 C.F.R. § 1.56.

75. On information and belief, Applicants and/or the prosecuting attorney were aware of the paper “Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web,” by David Karger et al. (the “Karger paper”) during prosecution of the ’222 application. Kove produced the Karger paper during discovery in this litigation, and, on information and belief, the Karger paper was in the same file since January 3, 1999 and during the prosecution of the ’222 application. Kove identified Mr. Overton as the custodian for this document.

76. The Karger paper is highly material prior art to the ’640 patent, but neither the prosecuting attorney nor the Applicants disclosed the Karger paper to the USPTO. More specifically, at least claims 17, 18 and 24 would not have issued had the Karger paper been before the USPTO during prosecution.

77. The Karger paper is material to the ’640 patent at least because it describes “caching protocols for distributed networks that can be used to decrease or eliminate the occurrence of hot spots” in networks such as the World Wide Web. The Karger paper is cited in AWS’s Final Invalidity Contentions, which are hereby incorporated by reference, as prior art that “either

anticipate[s] or render[s] obvious the patents-in-suit in combination with one or more other references and/or the knowledge of one of ordinary skill in the art around the time of the invention.”

78. On information and belief, Applicants and/or the prosecuting attorney made a deliberate decision to withhold the Karger paper from the USPTO with knowledge of the Karger paper’s materiality and with the specific intent to deceive the USPTO.

79. On information and belief, the ’640 patent is unenforceable because Applicants and/or the prosecuting attorney failed to disclose the Karger paper and made material misrepresentations about the prior art to the USPTO.

80. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the ’640 patent are unenforceable for the reasons set forth in paragraphs 69–79 above.

**TWELFTH COUNTERCLAIM: UNPATENTABILITY OF U.S. PATENT NO. 7,103,640**

81. The allegations of paragraphs 1 through 80 are realleged and reincorporated by reference as if fully set forth herein.

82. The ’640 patent does not claim patentable subject matter under 35 U.S.C. § 101 and is thus unpatentable.

83. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to unpatentability of the ’640 patent.

84. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the ’640 patent are unpatentable under 35 U.S.C. § 101.

**THIRTEENTH COUNTERCLAIM: DOUBLE-PATENTING**

85. The allegations of paragraphs 1 through 84 are realleged and reincorporated by reference as if fully set forth herein.

86. Claims 1, 3, 6, 10, 14, 17, 23, 24, 30 and 31 of the '978 patent are invalid under the doctrine of obviousness-type double patenting.

87. Therefore, there exists an actual controversy between AWS and Counter-Defendant related to the invalidity of claims of the '978 patent.

88. Pursuant to the Federal Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, AWS requests a declaration of the Court that the claims of the '978 patent are invalid under the doctrine of obviousness-type double patenting.

**PRAYER FOR RELIEF**

WHEREFORE, AWS seeks judgment in its favor and against Counter-Defendant declaring:

- A. That the claims of the '978 patent are not infringed by AWS;
- B. That the claims of the '978 patent are invalid;
- C. That the claims of the '978 patent are unenforceable;
- D. That the '978 patent is unpatentable;
- E. That the claims of the '170 patent are not infringed by AWS;
- F. That the claims of the '170 patent are invalid;
- G. That the claims of the '170 patent are unenforceable;
- H. That the '170 patent is unpatentable;
- I. That the claims of the '640 patent are not infringed by AWS;
- J. That the claims of the '640 patent are invalid;



- K. That the claims of the '640 patent are unenforceable;
- L. That the '640 patent is unpatentable;
- M. That the claims of the '978 patent are invalid under the doctrine of obviousness-type double patenting.
- N. That this is an exceptional case and AWS is entitled to attorneys' fees, costs, and expenses in this action; and
- O. That the Court may grant AWS such other and further relief as it may deem just and proper.

**JURY DEMAND**

AWS demands a trial by jury on all issues triable to a jury.

May 1, 2020

Respectfully submitted,

/s/ Terri L. Mascherin

Adam G. Unikowsky  
JENNER & BLOCK LLP  
1099 New York Ave. NW Suite 900  
Washington, DC 20001  
(202) 639-6000  
aunikowsky@jenner.com

Terri L. Mascherin  
Timothy J. Barron  
Michael T. Werner  
JENNER & BLOCK LLP  
353 N. Clark St.  
Chicago, IL 60654  
(312) 222-9350  
tmascherin@jenner.com  
tbarron@jenner.com  
mwerner@jenner.com

*Attorneys for Amazon Web Services, Inc.*

**CERTIFICATE OF SERVICE**

I hereby certify that on May 1, 2020, I caused a true and correct copy of the foregoing motion to be served on all counsel of record via the Court's ECF system

/s/ Michael T. Werner

Michael T. Werner  
Jenner & Block LLP  
353 N. Clark Street  
Chicago, Illinois 60654  
(312) 222-9350