

**IN THE UNITED STATES DISTRICT COURT  
FOR THE WESTERN DISTRICT OF TEXAS  
WACO DIVISION**

STREAMSCALE, INC.,	)	
	)	
Plaintiff,	)	
	)	
v.	)	Civil No. 6:21-cv-00198-ADA
	)	
CLOUDERA, INC.,	)	<b>JURY TRIAL DEMANDED</b>
AUTOMATIC DATA PROCESSING, INC.,	)	
EXPERIAN PLC, WARGAMING	)	
(AUSTIN), INC., and	)	
INTEL CORPORATION,	)	
	)	
Defendants.	)	

**FIRST AMENDED COMPLAINT FOR PATENT INFRINGEMENT**

Plaintiff StreamScale, Inc. (“Plaintiff” or “StreamScale”) files this First Amended Complaint for patent infringement against Defendants Cloudera, Inc. (“Cloudera”), Automatic Data Processing, Inc.<sup>1</sup> (“ADP”), Experian plc (“Experian”), Wargaming (Austin), Inc. (“Wargaming”), and Intel Corporation (“Intel”) (collectively, “Defendants”) alleging as follows:

**NATURE OF SUIT**

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

---

<sup>1</sup> Defendant Automatic Data Processing, Inc. indicated in its Answer that it would move to formally substitute ADP, Inc. in place of itself. Answer at 1 n.1, *StreamScale, Inc. v. Cloudera, Inc.*, No. 6:21-cv-00198-ADA (W.D. Tex. May 10, 2021), ECF No. 24. StreamScale does not anticipate opposing such a motion, but at this time, neither Automatic Data Processing, Inc. nor ADP, Inc. have filed such a motion. To maintain the status quo, StreamScale, Inc. has again named Automatic Data Processing, Inc. in this First Amended Complaint as it did in the Original Complaint. Original Complaint for Patent Infringement, *StreamScale, Inc. v. Cloudera, Inc.*, No. 6:21-cv-00198-ADA (W.D. Tex. Mar. 2, 2021), ECF No. 1.

2. StreamScale owns multiple patents relating to accelerated erasure coding. StreamScale's patented technology is a cornerstone of modern data storage, especially cloud-based data storage.

3. Data storage protection from loss used to be a matter of replicating the data. Data replication resulted in redundant data drives, and that redundancy provided an enhanced measure of data availability along with some measure of fault tolerance. For example, if one of the data drives were to be corrupted, the original data would still be available on a redundant disk.

4. Data replication is highly inefficient and no longer commercially practicable. Take a triple replication scheme for example. If a user desired to save some quantum of data, say 100 GB, it would require 300 GB of data storage to save that 100 GB of data. That is only a 33% utilization of storage capacity. And that measure of efficiency gets worse as the amount of redundancy in a system increases. Triple replication is also incredibly expensive because you need to buy three times the capacity of your original data. Triple replication further requires the additional, redundant capacity to be packaged, powered, and serviced.

5. Systems that employ accelerated erasure coding as patented by StreamScale enable scalable, high-performance data storage systems that can outperform other systems and do so at lower cost. StreamScale's inventions significantly reduce storage overhead while achieving similar or better fault tolerance than prior systems and methods, and are a quantum leap forward from prior systems.

6. At a high level, erasure coding uses specially designed systems to transform a block of original data to be stored into one or more blocks of encoded data that can be distributed across numerous storage devices or drives. The original data can be reconstructed from the encoded data, even if some portions of the original data are lost or unavailable. The data encoding and decoding

processes are time and energy intensive. If erasure coding is performed without appropriately configured computers using appropriately organized instructions, it can appear to have only limited practical applicability. Indeed, the widespread view in the industry before the work of StreamScale was that there was no way to employ erasure coding at high speeds, including so-called “cache line speeds.”

7. With its accelerated erasure coding technology, StreamScale achieved what was thought to be impossible. StreamScale achieved in one embodiment more than an order of magnitude performance increase in actual system performance. Rather than remaining an unobtainable goal with very limited application, storage systems based on StreamScale’s accelerated erasure coding immediately became practical and thus had newfound applicability to the data storage industry, among others.

8. The innovations described in—and protected by—StreamScale’s patents have been incorporated into products and services offered by Cloudera, ADP, Experian, and Wargaming. For its part, Intel has induced infringement by at least Cloudera, ADP, Experian, and Wargaming through Intel’s collaboration with Cloudera relating to accelerated erasure coding.

### **PARTIES**

9. Plaintiff StreamScale, Inc. is a corporation duly organized and existing under the laws of the State of Texas, having a principal place of business at 7215 Bosque Blvd., Suite 203, Waco, Texas 76710. StreamScale is the owner of record of the Patents-in-Suit in this action.

10. Defendant Cloudera, Inc. (“Cloudera”) is a corporation organized under the laws of the State of Delaware. Cloudera maintains an office in this judicial district at 515 Congress, Suite 1300, Austin, Texas 78701. Cloudera can be served with process through its registered agent in the State of Texas, Corporation Service Company d/b/a CSC – Lawyers Incorporating Service Company, 211 East 7th Street, Suite 620, Austin, Texas 78701-3218.

11. Defendant Automatic Data Processing, Inc. (“ADP”) is a corporation organized under the laws of the State of Delaware. ADP maintains offices in this judicial district, including at (i) 6500 River Place Blvd., Bldg VII, Austin, Texas 78730, (ii) 1851 North Resler, El Paso, Texas 79912, (iii) 7650 San Felipe Dr., El Paso, Texas 79912, and (iv) 211 North Loop 1604 East, San Antonio, Texas 78232. ADP can be served with process through its registered agent in the State of Texas, C T Corporation System, 1999 Bryan St., Suite 900, Dallas, Texas 75201.

12. Defendant Experian PLC is a public limited company registered and incorporated under the laws of the Bailiwick of Jersey, having a principal place of business at Newenham House, Northern Cross, Malahide Road, Dublin 17, D17 AY61, Ireland, and registered office at 22 Grenville Street, St Helier, Jersey JE4 8PX, Channel Islands.

13. Defendant Wargaming (Austin), Inc. (“Wargaming”) is a corporation organized under the laws of the State of Delaware, having a principal place of business at 11001 Lakeline Blvd., Austin, Texas 78717. Wargaming can be served with process through its registered agent in the State of Texas, C T Corporation System, 1999 Bryan St., Suite 900, Dallas, Texas 75201.

14. Defendant Intel Corporation (“Intel”) is a corporation organized under the laws of the State of Delaware. Intel maintains an office in this judicial district at 9442 N. Capital of Texas Hwy, Bldg 2, Suite 600, Austin, Texas 78759. Intel can be served with process through its registered agent in the State of Texas, C T Corporation System, 1999 Bryan St., Suite 900, Dallas, Texas 75201.

15. Collectively, Cloudera, ADP, Experian, Wargaming, and Intel are referred to herein as the “Defendants.”

### **JURISDICTION AND VENUE**

16. This action arises under the patent laws of the United States, 35 U.S.C. § 101, *et seq.* This Court has jurisdiction over this action pursuant to 28 U.S.C. §§ 1331 and 1338(a).

17. Cloudera is subject to personal jurisdiction in this Court. This Court has personal jurisdiction over Cloudera because Cloudera has engaged in continuous, systematic, and substantial activities within this State, including substantial marketing and sales of products and services within this State and this District. Furthermore, upon information and belief, this Court has personal jurisdiction over Cloudera because Cloudera has committed acts giving rise to StreamScale's claims for patent infringement within and directed to this District.

18. Upon information and belief, Cloudera has conducted and does conduct substantial business in this forum, directly and/or through subsidiaries, agents, representatives, or intermediaries, such substantial business including but not limited to: (i) at least a portion of the acts of infringement alleged herein; (ii) purposefully and voluntarily placing one or more infringing products and services into the stream of commerce with the expectation that they will be purchased by consumers in this forum; and/or (iii) regularly doing or soliciting business, engaging in other persistent courses of conduct, or deriving substantial revenue from goods and services provided to individuals in Texas and in this judicial district. Thus, Cloudera is subject to this Court's specific and general personal jurisdiction pursuant to due process and the Texas Long Arm Statute.

19. Upon information and belief, Cloudera has committed acts of infringement in this District and has one or more regular and established places of business within this District under 28 U.S.C. § 1400(b). Thus, venue is proper in this District under 28 U.S.C. § 1400(b).

20. Cloudera maintains a permanent physical presence within this District. For example, it maintains at least the office location at 515 Congress, Suite 1300, Austin, Texas 78701. Cloudera employs numerous employees who work at Cloudera's location(s) in this District.

21. Cloudera's location(s) in this District, including at least those identified in paragraph 20 above, are regular and established places of business under 28 U.S.C. § 1391, 28 U.S.C. § 1400(b), and *In re Cray, Inc.*, 871 F.3d 1355, 1360 (Fed. Cir. 2017).

a. Cloudera's location(s) in this District, including at least those identified in paragraph 20 above, are physical, geographical location(s) in this District. Each office location comprises one or more buildings or office spaces from which the business of Cloudera is carried out. The location(s) are set apart for the purpose of carrying out Cloudera's business, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services. Indeed, Cloudera itself advertises its physical location(s) in this District as places of its business.

b. Cloudera's location(s) in this District, including at least those identified in paragraph 20 above, are regular and established.

c. Cloudera's location(s) in this District, including at least those identified in paragraph 20 above, are places of business of Cloudera. Cloudera conducts business from its location(s) in this District, including at least those identified in paragraph 20 above, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services.

d. Cloudera's location(s) in this District, including at least those identified in paragraph 20 above, are physical, geographical location(s) in this District from which Cloudera carries out its business.

e. Cloudera employees work at Cloudera's location(s), including at least those identified in paragraph 20 above. Upon information and belief, these Cloudera employees are regularly and physically present at Cloudera's location(s), including at least those

identified in paragraph 20 above, during business hours and they are conducting Cloudera's business while working there.

22. ADP is subject to personal jurisdiction in this Court. This Court has personal jurisdiction over ADP because ADP has engaged in continuous, systematic, and substantial activities within this State, including substantial marketing and sales of products and services within this State and this District. Furthermore, upon information and belief, this Court has personal jurisdiction over ADP because ADP has committed acts giving rise to StreamScale's claims for patent infringement within and directed to this District.

23. Upon information and belief, ADP has conducted and does conduct substantial business in this forum, directly and/or through subsidiaries, agents, representatives, or intermediaries, such substantial business including but not limited to: (i) at least a portion of the acts of infringement alleged herein; (ii) purposefully and voluntarily placing one or more infringing products into the stream of commerce with the expectation that they will be purchased by consumers in this forum; and/or (iii) regularly doing or soliciting business, engaging in other persistent courses of conduct, or deriving substantial revenue from goods and services provided to individuals in Texas and in this judicial district. Thus, ADP is subject to this Court's specific and general personal jurisdiction pursuant to due process and the Texas Long Arm Statute.

24. Upon information and belief, ADP has committed acts of infringement in this District and has one or more regular and established places of business within this District under 28 U.S.C. § 1400(b). Thus, venue is proper in this District under 28 U.S.C. § 1400(b).

25. ADP maintains a permanent physical presence within this District. For example, it maintains office locations at (i) ADP Austin, 6500 River Place Blvd. Bldg. VII, Austin, Texas 78730; (ii) ADP El Paso, 1851 North Resler, El Paso, Texas 79912; (iii) ADP El Paso, 7650 San

Felipe Drive, El Paso, Texas 79912; and (iv) ADP San Antonio, 211 North Loop 1604 East, San Antonio, Texas 78232. ADP employs employees who work at ADP's locations in this District.

26. ADP's location(s) in this District, including at least those identified in paragraph 25 above, are regular and established places of business under 28 U.S.C. § 1391, 28 U.S.C. § 1400(b), and *In re Cray, Inc.*, 871 F.3d 1355, 1360 (Fed. Cir. 2017).

a. ADP's location(s) in this District, including at least those identified in paragraph 25 above, are physical, geographical location(s) in this District. Each office location comprises one or more buildings or office spaces from which the business of ADP is carried out. The location(s) are set apart for the purpose carrying out ADP's business, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services. Indeed, ADP itself advertises its physical location(s) in this District as places of its business, and it features commercial signage at many of these location(s).

b. ADP's location(s) in this District, including at least those identified in paragraph 25 above, are regular and established. ADP features commercial signage at many of the location(s) identifying the location as a regular and established place of ADP's business.

c. ADP's location(s) in this District, including at least those identified in paragraph 25 above, are places of business of ADP. ADP conducts business from its location(s) in this District, including at least those identified in paragraph 25 above, including but not limited to making, using selling, offering for sale, and/or supporting infringing products and services.



d. ADP's location(s) in this District, including at least those identified in paragraph 25 above, are physical, geographical location(s) in this District from which ADP carries out its business.

e. ADP employees work at ADP's location(s), including at least those identified in paragraph 25 above. Upon information and belief, these ADP employees are regularly and physically present at ADP's location(s), including at least those identified in paragraph 25 above, during business hours and they are conducting ADP's business while working there.

27. Experian is subject to personal jurisdiction in this Court. This Court has personal jurisdiction over Experian because, upon information and belief, Experian has engaged in continuous, systematic, and substantial activities within this State, for example with and through its corporate subsidiaries CSIdentity Corporation and Experian Information Solutions, Inc. Upon information and belief, Experian's continuous, systematic, and substantial activities within this State include substantial marketing and sales of products and services within this State and this District, including for example through Experian's corporate subsidiaries CSIdentity Corporation and Experian Information Solutions, Inc. Furthermore, upon information and belief, this Court has personal jurisdiction over Experian because Experian has committed acts giving rise to StreamScale's claims for patent infringement within and directed to this District.

28. Upon information and belief, Experian has conducted and does conduct substantial business in this forum, directly and/or through subsidiaries, agents, representatives, or intermediaries, such substantial business including but not limited to: (i) at least a portion of the acts of infringement alleged herein; (ii) purposefully and voluntarily placing one or more infringing products into the stream of commerce with the expectation that they will be purchased

by consumers in this forum; and/or (iii) regularly doing or soliciting business, engaging in other persistent courses of conduct, or deriving substantial revenue from goods and services provided to individuals in Texas and in this judicial district. Thus, Experian is subject to this Court's specific and general personal jurisdiction pursuant to due process and the Texas Long Arm Statute.

29. To the extent Experian is not subject to jurisdiction in any State's courts of general jurisdiction, this Court has personal jurisdiction of Experian pursuant to Federal Rule of Civil Procedure 4(k)(2) because StreamScale's claims arise under federal law and exercising jurisdiction is consistent with the United States Constitution and laws.

30. Upon information and belief, Experian has committed acts of infringement in this District and has, itself or through its corporate subsidiaries, one or more regular and established places of business within this District under 28 U.S.C. § 1400(b). Thus, venue is proper in this District under 28 U.S.C. § 1400(b).

31. Experian, including for example through Experian's corporate subsidiaries CSIdentity Corporation and Experian Information Solutions, Inc., maintains a permanent physical presence within this District. For example, it maintains at least the office location at 1501 South MoPac Expressway, Austin, Texas 78746. Experian employs employees who work at Experian's location(s) in this District.

32. Experian's location(s) in this District, including at least those identified in paragraph 31 above, are regular and established places of business under 28 U.S.C. § 1391, 28 U.S.C. § 1400(b), and *In re Cray, Inc.*, 871 F.3d 1355, 1360 (Fed. Cir. 2017).

a. Experian's location(s) in this District, including for example those identified in paragraph 31 above, are physical, geographical location(s) in this District. Each office location comprises one or more buildings or office spaces from which the

business of Experian is carried out. The location(s) are set apart for the purpose of carrying out Experian's business, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services. Indeed, Experian itself advertises its physical location(s) in this District as places of its business, and it features commercial signage at these location(s).

b. Experian's location(s) in this District, including at least those identified in paragraph 31 above, are regular and established. Experian features commercial signage at the location(s) identifying the location as a regular and established place of Experian's business.

c. Experian's location(s) in this District, including at least those identified in paragraph 31 above, are places of business of Experian, including at least Experian's corporate subsidiaries CSIdentity Corporation and Experian Information Solutions, Inc. Experian conducts business from its location(s) in this District, including at least those identified in paragraph 31 above, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services.

d. Experian's location(s) in this District, including at least those identified in paragraph 31 above, are physical, geographical location(s) in this District from which Experian carries out its business.

e. Experian employees work at Experian's location(s), including at least those identified in paragraph 31 above. Upon information and belief, these Experian employees are regularly and physically present at Experian's location(s), including at least those identified in paragraph 31 above, during business hours and they are conducting Experian's business while working there.

33. To the extent Experian is found not reside in the United States, venue is nonetheless proper in this Court as to Experian pursuant to 28 U.S.C. § 1391(c)(3).

34. Wargaming is subject to personal jurisdiction in this Court. This Court has personal jurisdiction over Wargaming because Wargaming has engaged in continuous, systematic, and substantial activities within this State, including substantial marketing and sales of products and services within this State and this District. Furthermore, upon information and belief, this Court has personal jurisdiction over Wargaming because Wargaming has committed acts giving rise to StreamScale's claims for patent infringement within and directed to this District.

35. Upon information and belief, Wargaming has conducted and does conduct substantial business in this forum, directly and/or through subsidiaries, agents, representatives, or intermediaries, such substantial business including but not limited to: (i) at least a portion of the acts of infringement alleged herein; (ii) purposefully and voluntarily placing one or more infringing products and services into the stream of commerce with the expectation that they will be purchased by consumers in this forum; and/or (iii) regularly doing or soliciting business, engaging in other persistent courses of conduct, or deriving substantial revenue from goods and services provided to individuals in Texas and in this judicial district. Thus, Wargaming is subject to this Court's specific and general personal jurisdiction pursuant to due process and the Texas Long Arm Statute.

36. Upon information and belief, Wargaming has committed acts of infringement in this District and has one or more regular and established places of business within this District under 28 U.S.C. § 1400(b). Thus, venue is proper in this District under 28 U.S.C. § 1400(b).

37. Wargaming maintains a permanent physical presence within this District. For example, it maintains at least the office location at 11001 Lakeline Blvd., Austin, Texas 78717. Wargaming employs numerous employees who work at Wargaming's location(s) in this District.

38. Wargaming's location(s) in this District, including at least those identified in paragraph 37 above, are regular and established places of business under 28 U.S.C. § 1391, 28 U.S.C. § 1400(b), and *In re Cray, Inc.*, 871 F.3d 1355, 1360 (Fed. Cir. 2017).

a. Wargaming's location(s) in this District, including at least those identified in paragraph 37 above, are physical, geographical location(s) in this District. Each office location comprises one or more buildings or office spaces from which the business of Wargaming is carried out. The location(s) are set apart for the purpose of carrying out Wargaming's business, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services. Indeed, Wargaming itself advertises its physical location(s) in this District as places of its business.

b. Wargaming's location(s) in this District, including at least those identified in paragraph 37 above, are regular and established.

c. Wargaming's location(s) in this District, including at least those identified in paragraph 37 above, are places of business of Wargaming. Wargaming conducts business from its location(s) in this District, including at least those identified in paragraph 37 above, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services.

d. Wargaming's location(s) in this District, including at least those identified in paragraph 37 above, are physical, geographical location(s) in this District from which Wargaming carries out its business.

e. Wargaming employees work at Wargaming's location(s), including at least those identified in paragraph 37 above. Upon information and belief, these Wargaming employees are regularly and physically present at Wargaming's location(s), including at least those identified in paragraph 37 above, during business hours and they are conducting Wargaming's business while working there.

39. Intel is subject to personal jurisdiction in this Court. This Court has personal jurisdiction over Intel because Intel has engaged in continuous, systematic, and substantial activities within this State, including substantial marketing and sales of products and services within this State and this District. Furthermore, upon information and belief, this Court has personal jurisdiction over Intel because Intel has committed acts giving rise to StreamScale's claims for patent infringement within and directed to this District.

40. Upon information and belief, Intel has conducted and does conduct substantial business in this forum, directly and/or through subsidiaries, agents, representatives, or intermediaries, such substantial business including but not limited to: (i) at least a portion of the acts of infringement alleged herein; (ii) purposefully and voluntarily placing one or more infringing products into the stream of commerce with the expectation that they will be purchased by consumers in this forum; and/or (iii) regularly doing or soliciting business, engaging in other persistent courses of conduct, or deriving substantial revenue from goods and services provided to individuals in Texas and in this judicial district. Thus, Intel is subject to this Court's specific and general personal jurisdiction pursuant to due process and the Texas Long Arm Statute.

41. Upon information and belief, Intel has committed acts of infringement in this District and has one or more regular and established places of business within this District under 28 U.S.C. § 1400(b). Thus, venue is proper in this District under 28 U.S.C. § 1400(b).

42. Intel maintains a permanent physical presence within this District. For example, it maintains at least the office location at 9442 N. Capital of Texas Hwy, Bldg 2, Suite 600, Austin, Texas 78759. Intel employs numerous employees who work at Intel's location(s) in this District.

43. Intel's location(s) in this District, including at least those identified in paragraph 42 above, are regular and established places of business under 28 U.S.C. § 1391, 28 U.S.C. § 1400(b), and *In re Cray, Inc.*, 871 F.3d 1355, 1360 (Fed. Cir. 2017).

a. Intel's location(s) in this District, including at least those identified in paragraph 42 above, are physical, geographical location(s) in this District. Each office location comprises one or more buildings or office spaces from which the business of Intel is carried out. The location(s) are set apart for the purpose of carrying out Intel's business, including but not limited to the acts of infringement alleged herein. Indeed, Intel itself advertises its physical location(s) in this District as places of its business.

b. Intel's location(s) in this District, including at least those identified in paragraph 42 above, are regular and established.

c. Intel's location(s) in this District, including at least those identified in paragraph 42 above, are places of business of Intel. Intel conducts business from its location(s) in this District, including at least those identified in paragraph 42 above, including but not limited to, making, using, selling, offering for sale, and/or supporting infringing products and services.

d. Intel's location(s) in this District, including at least those identified in paragraph 42 above, are physical, geographical location(s) in this District from which Intel carries out its business.

e. Intel employees work at Intel's location(s), including at least those identified in paragraph 42 above. Upon information and belief, these Intel employees are regularly and physically present at Intel's location(s), including at least those identified in paragraph 42 above, during business hours and they are conducting Intel's business while working there.

### **FACTUAL ALLEGATIONS**

#### **I. PATENTS-IN-SUIT**

44. U.S. Patent No. 8,683,296 ("the '8-296 Patent") is entitled "Accelerated Erasure Coding System and Method." The '8-296 Patent duly and legally issued on March 25, 2014, from U.S. Patent Application No. 13/341,833, filed on December 30, 2011. StreamScale is the current owner of all rights, title, and interest in and to the '8-296 Patent. A true and correct copy of the '8-296 Patent is attached hereto as Exhibit A and is incorporated by reference herein.

45. U.S. Patent No. 9,160,374 ("the '374 Patent") is entitled "Accelerated Erasure Coding System and Method." The '374 Patent duly and legally issued on October 13, 2015, from U.S. Patent Application No. 14/223,740, filed on March 24, 2014. The '374 Patent is a continuation of U.S. Patent Application No. 13/341,833, filed on December 30, 2011, now U.S. Patent No. 8,683,296. The '374 Patent is entitled to the benefit of the December 30, 2011 filing date of application No. 13/341,833. StreamScale is the current owner of all rights, title, and interest in and to the '374 Patent. A true and correct copy of the '374 Patent is attached hereto as Exhibit B and is incorporated by reference herein.

46. U.S. Patent No. 9,385,759 ("the '759 Patent") is entitled "Accelerated Erasure Coding System and Method." The '759 Patent duly and legally issued on July 5, 2016, from U.S. Patent Application No. 14/852,438, filed on September 11, 2015. The '759 Patent is a continuation of U.S. Patent Application No. 14/223,740, filed on March 24, 2014, now U.S. Patent



No. 9,160,374. U.S. Patent No. 9,160,374 is a continuation of U.S. Patent Application No. 13/341,833, filed on December 30, 2011, now U.S. Patent No. 8,683,296. The '759 Patent is entitled to the benefit of the December 30, 2011 filing date of application No. 13/341,833. StreamScale is the current owner of all rights, title, and interest in and to the '759 Patent. A true and correct copy of the '759 Patent is attached hereto as Exhibit C and is incorporated by reference herein.

47. U.S. Patent No. 10,003,358 (“the '358 Patent”) is entitled “Accelerated Erasure Coding System and Method.” The '358 Patent duly and legally issued on June 19, 2018, from U.S. Patent Application No. 15/201,196, filed on July 1, 2016. The '358 Patent is a continuation of U.S. Patent Application No. 14/852,438, filed on September 11, 2015, now U.S. Patent No. 9,385,759. U.S. Patent No. 9,385,759 is a continuation of U.S. Patent Application No. 14/223,740, filed on March 24, 2014, now U.S. Patent No. 9,160,374. U.S. Patent No. 9,160,374 is a continuation of U.S. Patent Application No. 13/341,833, filed on December 30, 2011, now U.S. Patent No. 8,683,296. The '358 Patent is entitled to the benefit of the December 30, 2011 filing date of application No. 13/341,833. StreamScale is the current owner of all rights, title, and interest in and to the '358 Patent. A true and correct copy of the '358 Patent is attached hereto as Exhibit D and is incorporated by reference herein. On or about February 23, 2021, StreamScale filed a Petition for Correction of Inventorship Under 37 C.F.R. § 1.324, including associated documentation and fees, with the United States Patent and Trademark Office requesting the correction of inventorship of the '358 Patent to include inventor Sarah Mann, who was not named as an inventor through error. True and correct copies of that Petition and associated documentation are attached as Exhibit E, and that material is incorporated by reference herein.

48. U.S. Patent No. 10,291,259 (“the ’259 Patent”) is entitled “Accelerated Erasure Coding System and Method.” The ’259 Patent duly and legally issued on May 14, 2019, from U.S. Patent Application No. 15/976,175, filed on May 10, 2018. The ’259 Patent is a continuation of U.S. Patent Application No. 15/201,196, filed on July 1, 2016, now U.S. Patent No. 10,003,358. U.S. Patent No. 10,003,358 is a continuation of U.S. Patent Application No. 14/852,438, filed on September 11, 2015, now U.S. Patent No. 9,385,759. U.S. Patent No. 9,385,759 is a continuation of U.S. Patent Application No. 14/223,740, filed on March 24, 2014, now U.S. Patent No. 9,160,374. U.S. Patent No. 9,160,374 is a continuation of U.S. Patent Application No. 13/341,833, filed on December 30, 2011, now U.S. Patent No. 8,683,296. The ’259 Patent is entitled to the benefit of the December 30, 2011 filing date of application No. 13/341,833. StreamScale is the current owner of all rights, title, and interest in and to the ’259 Patent. A true and correct copy of the ’259 Patent is attached hereto as Exhibit F and is incorporated by reference herein. On or about February 23, 2021, StreamScale filed a Petition for Correction of Inventorship Under 37 C.F.R. § 1.324, including associated documentation and fees, with the United States Patent and Trademark Office requesting the correction of inventorship of the ’259 Patent to include inventor Sarah Mann, who was not named as an inventor through error. True and correct copies of that Petition and associated documentation are attached as Exhibit G, and that material is incorporated by reference herein.

49. U.S. Patent No. 10,666,296 (“the ’10-296 Patent”) is entitled “Accelerated Erasure Coding System and Method.” The ’10-296 Patent duly and legally issued on May 26, 2020, from U.S. Patent Application No. 16/358,602, filed on March 19, 2019. The ’10-296 Patent is a continuation of U.S. Patent Application No. 15/976,175, filed on May 10, 2018, now U.S. Patent No. 10,291,259. U.S. Patent No. 10,291,259 is a continuation of U.S. Patent Application

No. 15/201,196, filed on July 1, 2016, now U.S. Patent No. 10,003,358. U.S. Patent No. 10,003,358 is a continuation of U.S. Patent Application No. 14/852,438, filed on September 11, 2015, now U.S. Patent No. 9,385,759. U.S. Patent No. 9,385,759 is a continuation of U.S. Patent Application No. 14/223,740, filed on March 24, 2014, now U.S. Patent No. 9,160,374. U.S. Patent No. 9,160,374 is a continuation of U.S. Patent Application No. 13/341,833, filed on December 30, 2011, now U.S. Patent No. 8,683,296. The '10-296 Patent is entitled to the benefit of the December 30, 2011 filing date of application No. 13/341,833. StreamScale is the current owner of all rights, title, and interest in and to the '10-296 Patent. A true and correct copy of the '10-296 Patent is attached hereto as Exhibit H and is incorporated by reference herein.

50. Collectively, the '8-296 Patent, the '374 Patent, the '759 Patent, the '358 Patent, the '259 Patent, and the '10-296 Patent are referred to herein as the “Patents-in-Suit.”

## **II. ACCELERATED ERASURE CODING INFRINGEMENT**

51. As further discussed below, Cloudera, ADP, Experian, and Wargaming (the “EC System Defendants”) directly and/or indirectly infringed—and continue to directly and/or indirectly infringe—each of the Patents-in-Suit by engaging in acts constituting infringement under 35 U.S.C. § 271(a) and (b), including without limitation by one or more of making, using, selling, and/or offering to sell, in this District and elsewhere in the United States, and/or importing into this District and elsewhere in the United States, systems that incorporate Cloudera Erasure Coding Components. Cloudera Erasure Coding Components include Cloudera Distribution Including Apache Hadoop (“Cloudera CDH”), which may include any related components, and any Cloudera product or service that is substantially or reasonably similar, including but not limited to Cloudera Enterprise. The infringing systems that Cloudera runs that use the Cloudera Erasure Coding Components are the “Cloudera Infringing Products and Services.”

52. Systems built by Cloudera, ADP, Experian, and Wargaming with Cloudera CDH or substantially similar technology include accelerated erasure coding (“EC”) technology are the “EC Systems.” These Defendants are “EC System Defendants.”

53. Under typical configurations, the EC Systems that use the patented technology reduce storage cost by at least about 50% compared with triple replication. Upon information and belief, Cloudera and its collaborators recognized that accelerated erasure coding can greatly reduce storage overhead without sacrificing data reliability, which makes erasure coding a quite attractive alternative for big data storage, particularly for cold data.

54. EC technology is packaged and shipped with Cloudera CDH. Additionally, this EC technology is enabled by default in Cloudera CDH.

55. ADP has directly infringed, and continues to directly infringe, each of the Patents-in-Suit by engaging in acts constituting infringement under 35 U.S.C. § 271(a), including without limitation by one or more of making, using, selling and/or offering to sell, in this District and elsewhere in the United States, and/or importing into this District and elsewhere in the United States, at least ADP’s products and services that use and/or incorporate the Cloudera Erasure Coding Components, including but not limited to DataCloud, and any ADP product or service that is substantially or reasonably similar (the “ADP Infringing Products and Services”).

56. As its name implies, data is core to ADP’s business. Upon information and belief, ADP, a provider of human capital management solutions, is responsible for getting one in six Americans paid today, which puts tremendous data in ADP’s hands.

57. Upon information and belief, ADP is now putting that data to use and generating a new revenue stream. For example, upon information and belief, ADP has built at least a product called DataCloud, which employs Cloudera Erasure Coding Components, that aggregates

information across ADP's 600,000 clients and generates insights to help clients prevent employee churn, ensure salary equality, and maximize human resources.

58. Upon information and belief, ADP was able to use DataCloud to identify the top one percent of at-risk employees in a pilot account, and learned that within that group, turnover was actually 50 percent. When removing that top one percent from the overall analysis, average turnover dropped to nine percent. DataCloud helped the client focus on a small population of at-risk employees where they could make a meaningful impact that would drastically improve the company's overall churn; without this insight, they would have spread retention efforts across the employee base, requiring more time and resources with a less targeted approach and having a lower impact overall.

59. Reducing employee churn has far-reaching business impacts. The cost of losing one employee is more than a simple hiring replacement. Recruiting and interviewing for that person's replacement is costly. Productivity is lost while the new hire gets up to speed. Risk of others on the team leaving increases when they're forced to pick up the slack. It's a ripple effect.

60. The value DataCloud offers is evidenced by the massive growth ADP has seen throughout its client base, driving greater success for ADP via this new revenue channel.

61. Upon information and belief, DataCloud stemmed from a strategic shift at ADP to move from primarily processing transactions to also providing insights based on its greatest asset: data. Upon considering building this product, ADP reached out to clients to gauge their interest in gaining insights based on aggregated and anonymized benchmarks developed from the data spanning ADP's customer base. But making the vision a reality presented a technological challenge. Upon information and belief, ADP's data was spread across data centers and

applications. It needed to be brought together for processing, exploration, and analysis. It wouldn't be feasible using traditional relational database technology.

62. ADP built DataCloud to allow for the storage and processing of large amounts of data in new ways. According to Jim Haas, Principal Architect of DataCloud, advanced data storage and prioritization technologies let companies "maximize CPU time and memory used," which for HR leaders means "getting the big tasks done faster."<sup>2</sup> KPMG reports that 42% of organizations will replace their existing HR software with a cloud-based solution, with most citing better functionality and higher business value as the motivation.<sup>3</sup> But the sheer amount of employee data, devices, access permissions, and historical data needed to effectively track current conditions and develop long-term policies can easily overwhelm standard infrastructure.

63. Upon information and belief, DataCloud employs Cloudera Enterprise, comprising a 200-terabyte (TB) lab and two 400-TB production data centers, each with replication for disaster recovery.

64. Upon information and belief, ten data domains feed DataCloud a billion records every quarter, including: (1) 600,000-plus client databases capturing information on 29 million people; (2) mainframe-based data from the 30 to 35 million pay cycles ADP executes annually, including compensation, time card punches, bonuses, overtime, and salary increases; (3) Oracle-based data from the 15 million HR functions managed by ADP annually, such as benefit deductions and elections, performance scores, and recruiting processes; (4) data from 15 other ADP

---

<sup>2</sup> Doug Bonderud, *HR Cloud Solutions: A Foundation for Better Decision Making*, ADP, <https://www.adp.com/spark/articles/2018/01/hr-cloud-solutions-a-foundation-for-better-decision-making.aspx> (last visited Jan. 19, 2021).

<sup>3</sup> See, e.g., 2016 HR Transformation Survey: Summary Report, KPMG, <https://assets.kpmg/content/dam/kpmg/in/pdf/2016/11/HR-Transformation-Survey-Summaryreport.pdf> (last visited Jan. 19, 2021).

departments—such as Marketing, Sales, Implementations, and Service—who leverage the platform as their enterprise data hub (EDH) so they may build their own data products; and (5) client data sets such as point-of-sale transactions and revenues.

65. DataCloud conforms job title and role categorizations across 600,000 companies into a comparable standard from which 500 billion aggregates are created. Those aggregates are used to build the benchmarks that are delivered to clients. Upon information and belief, Jim Haas, Principal Architect at ADP has explained, “the data is drawing everybody together . . . . Sometimes I call it ‘the little cluster that can’ because it’s just amazing what goes on in there in a day.”

66. Upon information and belief, the ADP Infringing Products and Services are configured to support accelerated erasure coding.

67. Experian has directly infringed, and continues to directly infringe, each of the Patents-in-Suit by engaging in acts constituting infringement under 35 U.S.C. §§ 271(a), including without limitation by one or more of making, using, selling and/or offering to sell, in this District and elsewhere in the United States, and/or importing into this District and elsewhere in the United States, at least Experian’s products and services that use and/or incorporate the Cloudera Infringing Products and Services, including but not limited to Experian Analytical Sandbox and Velcro, and any Experian product or service that is substantially or reasonably similar (the “Experian Infringing Products and Services”).

68. Upon information and belief, Experian integrated Cloudera Enterprise onto its cloud environment for its Credit Information Services, Decision Analytics, and Business Information Services business lines. Upon information and belief, Experian employs Cloudera Erasure Coding Components in Experian’s Ascend Technology Platform and Analytical Sandbox.

69. Experian is doing business in the United States and more particularly in this District, including at least through Experian's corporate subsidiaries CSIdentity Corporation and Experian Information Solutions, Inc., by making, using, selling, importing, and/or offering for sale the product and services that infringe one or more of the patent claims involved in this action.

70. Upon information and belief, with 15,000+ employees and annual revenues exceeding \$4 billion (USD), Experian is a global leader in credit reporting and marketing services that is comprised of four main business units: Credit Information Services, Decision Analytics, Business Information Services, and Marketing Services.

71. Experian Marketing Services ("EMS"), for example, helps marketers connect with customers through relevant communications across a variety of channels, driven by advanced analytics on an extensive database of geographic, demographic, and lifestyle data.

72. EMS has built its business on the effective collection, analysis, and use of data. Upon information and belief, as EMS's former VP of product strategy Jeff Hassemer once explained, "Experian has handled large amounts of data for a very long time: who consumers are, how they're connected, how they interact. We've done this over billions and quadrillions of records over time. But with the proliferation of channels and information that are now flowing into client organizations—social media likes, web interactions, email responses—that data has gotten so large that it's maxed the capacity of older systems. We needed to leap forward in our processing ability. We wanted to process data orders of magnitude faster so we could react to tomorrow's consumer."

73. Today's consumers leave a digital trail of behaviors and preferences for marketers to leverage so they can enhance the customer experience, and upon information and belief,



Experian's clients started asking for more frequent updates on consumers' latest purchasing behaviors, online browsing patterns, and social media activity so they can respond in real time.

74. Upon information and belief, Experian recognized that the data exhaust from these digital channels is massive and requires a technological infrastructure that can accommodate rapid processing, large-scale storage, and flexible analysis of multi-structured data. Experian's mainframes were hitting the tipping point in terms of performance, flexibility, and scalability. Given the need for immediacy of information and customization of data in real time for clients, EMS set an internal goal to process more than 100 million records of data per hour (28,000 records per second).

75. Upon information and belief, instead of trying to fit a square peg in a round hole, Experian went out and decided to build an architecture that could handle the new volumes of data that it manages and built a system that employs Cloudera CDH.

76. Upon information and belief, the Experian Infringing Products and Services are configured to support accelerated erasure coding.

77. Wargaming has directly infringed, and continues to directly infringe, each of the Patents-in-Suit by engaging in acts constituting infringement under 35 U.S.C. § 271(a), including without limitation by one or more of making, using, selling and/or offering to sell, in this District and elsewhere in the United States, and/or importing into this District and elsewhere in the United States, at least Wargaming's products and services that use and/or incorporate Cloudera Erasure Coding Components, including but not limited to Wargaming's Player Relationship Management Platform ("PRMP") in support of Wargaming's online games and massively multiplayer online ("MMO") games, and any Wargaming product or service that is substantially or reasonably similar (the "Wargaming Infringing Products and Services").

78. Wargaming provides strategic intelligence analytics services and coordinates the data services architecture for Wargaming MMO games. Wargaming is a global services hub for games developed, at least in part, by Wargaming Group Limited and accessible via the portal at [www.wargaming.net](http://www.wargaming.net). Wargaming provides data-driven insights, analysis, and reporting of wargaming.net projects, strategic planning, and global game design services through business analytics, production, central technology, and regional administrative departments.

79. Furthermore, Wargaming conducts general research on topics such as the gaming industry, player behavior, and game defects.

80. Wargaming serves more than 150 million registered players in its MMO games. Those millions of players generate massive amounts of data. Wargaming processes over 3 TB of data daily. Upon information and belief, it does this using systems built with Cloudera Erasure Coding Components.

81. Wargaming employees administrate and optimize a series of development and production clusters that employ Cloudera Erasure Coding Components.

82. Upon information and belief, the Wargaming Infringing Products and Services are configured to support accelerated erasure coding.

### **III. WIDESPREAD KNOWLEDGE OF STREAMSCALE'S PATENTS**

83. The United States Patent and Trademark Office published the patent application that ultimately led to the '8-296 Patent on July 4, 2013. The very next day, July 5, 2013, StreamScale sent a letter to USENIX, a computing systems association, notifying USENIX of StreamScale's pending patent applications and providing USENIX with advance notice of StreamScale's intent to issue a press release that StreamScale's then-patent-pending technology.<sup>4</sup>

---

<sup>4</sup> Exhibit I, Letter from Michael S. Adler, Counsel for StreamScale, to USENIX (July 5, 2013).

Upon information and belief, others in the industry, including but not limited to Intel, learned of StreamScale and its patent applications as a consequence of the letter StreamScale wrote to USENIX.

84. On July 10, 2013, while StreamScale awaited a response from USENIX, Intel publicly announced its excitement to support development of erasure code solutions. Intel explained that erasure codes reduce the size of data on disk by up to half versus traditional replication, decreases costs by more than 50%, and reduces both hardware requirement costs and power and cooling costs.<sup>5</sup> Intel explained that erasure code was a long overdue technology and Intel was excited to support, promote, and use it in cloud environments.<sup>6</sup>

85. On July 23, 2013, StreamScale issued a press release noting that its technology is protected by then-pending patent applications and was not “open source.”<sup>7</sup>

86. Having received the July 5, 2013 letter that StreamScale sent to USENIX, and following consultation with its attorneys, USENIX chose to comply with StreamScale’s request to remove certain papers and materials from its web site.

87. On or about August 3, 2013, individuals began posting missives online regarding StreamScale and its patent portfolio. Upon information and belief, H. Peter Anvin, an Intel employee was aware of at least some of these online postings. Indeed, upon information and belief,

---

<sup>5</sup> Exhibit J, Joe Arnold, Save Space: The Final Frontier—Erasure Codes with OpenStack Swift (July 10, 2013), *previously available at* <https://swiftstack.com/blog/2013/07/10/erasure-codes-with-openstack-swift/>.

<sup>6</sup> *Id.*

<sup>7</sup> Exhibit K, StreamScale Provides Notice of Ownership of Fastest Erasure Code Technology Disclosed at Fast ’13 (July 23, 2013).

at least Mr. Anvin commented on at least some of these online postings, including but not limited to on or about August 9, 2013.

88. Upon information and belief, in early March 2014, Intel employees again learned about StreamScale, its patented and patent-pending technology, and its relationship to ISA-L. On March 10, 2014, upon information and belief, one or more Intel employees reviewed and collected a significant quantity of information about StreamScale, its attorneys, and its patent applications. Upon information and belief, one or more Intel employees visited a number of specific pages on StreamScale's website, including (i) those detailing StreamScale's then-pending-patent applications, (ii) those summarizing StreamScale's company history and technology, (iii) those making recent new and press releases available to the public, (iv) those identifying StreamScale's employees and attorneys, and (v) those hosting academic papers authored by StreamScale's employees. Furthermore, upon information and belief, one or more Intel employees accessed and downloaded electronic copies of one or more of StreamScale's patents and patent applications, at least from StreamScale's website.

89. Indeed, upon information and belief, Intel was contacted in February or March 2014 and knew about potential issues involving StreamScale, StreamScale's patent-pending technology, and ISA-L. In August and September 2014, outside counsel for Intel corresponded with then-litigation counsel for StreamScale regarding a third party subpoena StreamScale issued to Intel involving StreamScale's intellectual property rights. Thus, upon information and belief, by mid-to-late September 2014, Intel had knowledge of StreamScale, StreamScale's issued and pending patents and intellectual property rights, and their relevance to ISA-L.

90. On September 19, 2014, StreamScale representatives spoke with legal counsel for Intel regarding StreamScale and StreamScale's patents. From September 19, 2014 through at least July 30, 2015, StreamScale and Intel engaged in extensive written and oral communications regarding StreamScale and StreamScale's patents. Through those discussions, StreamScale provided Intel with express notice of the '8-296 Patent and further provided Intel with express notice that the application that ultimately issued as the '374 Patent was recently allowed by the United States Patent & Trademark Office. Specifically, on or about July 20, 2015, Michael O'Shea, counsel for StreamScale, notified Robert A. Diehl, counsel for Intel, that StreamScale's patent portfolio included, among other assets, U.S. Patent No. 8,683,296, "Accelerated Erasure Coding System and Method" and U.S. Patent Application No. 14/223,740, "Accelerated Erasure Coding System and Method," a continuation of the '8-296 Patent that was "recently allowed." On or about July 20, 2015, Mr. O'Shea also provided Mr. Diehl with a claim chart describing Intel's infringement, including through ISA-L, of claim 1 of the '8-296 Patent.

91. Separately, on or about March 5, 2015, Tushar Gohad, an Intel employee, indicated that Jerasure and GF-complete were strategically important. Specifically, Mr. Gohad requested that Jerasure and GF-complete be backported to an earlier version of software. By that time, one of the authors of GF-Complete had publicly stated that StreamScale asserts that the use of GF-Complete (particularly as part of Jerasure 2.0 or later) or any similar software, method or code for erasure coding infringes StreamScale's issued United States Patent No. 8.683,296.

92. On or about April 29, 2015, counsel for StreamScale wrote on an online technical board and asked that the Jerasure 2.0 and GF-Complete libraries that had been republished be

removed.<sup>8</sup> The next day, April 30, 2015, upon information and belief, StreamScale’s post and a Techdirt article regarding StreamScale’s patent rights were brought to the attention of Paul Luse, another Intel employee, who responded “we are all well aware of the info you passed on :)” Upon information and belief, Mr. Luse then encouraged others to ignore StreamScale, indicating that was always the best option.

93. Since at least March 5, 2021, when Intel was served through its registered agent with the Original Complaint for Patent Infringement in this action, Return of Service, *StreamScale, Inc. v. Cloudera, Inc.*, No. 6:21-cv-00198-ADA (W.D. Tex. Mar. 10, 2021), ECF No. 14, Intel has had express knowledge of each of the Patents-in-Suit and its infringement thereof. Intel continues to actively induce infringement of the StreamScale patents.

#### **IV. INTEL’S INFRINGEMENT**

94. Upon information and belief, Intel is a Fortune 50 company, with revenues exceeding \$70 billion annually.

95. Intel has a long history with United States patent litigation. Upon information and belief, it employs several attorneys and counsel to manage its offensive and defensive patent litigation docket. In addition, upon information and belief, Intel employs several attorneys to evaluate, manage, and track patent assertions in its industry.

96. In addition, upon information and belief, Intel is a member or client of RPX Corporation (“RPX”).

97. RPX offers patent risk management services, including defensive patent buying, acquisition syndication, patent intelligence, and advisory services to its members and clients.

---

<sup>8</sup> Exhibit L, Michael A. O’Shea, counsel for StreamScale, post to Ubuntu entitled “StreamScale” (Apr. 29, 2015), available at <https://lists.ubuntu.com/archives/technical-board/2015-April/002100.html> (last visited May 24, 2021).

98. Also, upon information and belief, Intel is a member or client of Allied Security Trust (“AST”).

99. AST offers patent risk mitigation services to some of the world’s biggest technology companies and was created to combat unwanted patent assertions and litigation.

100. Intel has a publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals. “Intel’s own engineers concede that they avoid reviewing other, non-Intel patents so as to avoid willfully infringing them.” *Intel Corp. v. Future Link Sys., LLC*, 268 F. Supp. 3d 605, 623 (D. Del. 2017). Upon information and belief, Intel’s policy is designed to avoid possible triple damages for willful infringement.

101. In fact, upon information and belief, Intel has reprimanded its employees for inquiring about others’ intellectual property rights, including StreamScale’s patents. In early 2014, upon information and belief, Intel reprimanded one of its cloud software architects specifically for suggesting that there was a potential issue with ISA-L in connection with StreamScale.

102. Upon information and belief, Intel has rendered itself willfully blind to StreamScale’s patents and the intellectual property rights of others.

103. As further discussed below, Intel has indirectly infringed, and continues to indirectly infringe, each of the Patents-in-Suit by engaging in acts constituting infringement under 35 U.S.C. § 271(b), including without limitation by actively inducing infringement by the EC System Defendants through the deployment and/or use of Intel’s Intelligent Storage Acceleration Library (“ISA-L”).

104. Upon information and belief, Intel has been aware of the existence of the Patents-in-Suit beginning at least in 2014. In 2014, StreamScale provided Intel with express notice of the ’8-296 Patent. Upon information and belief, Intel knew that continuation patent applications

claimed priority to the '8-296 Patent. All of the Patents-in-Suit are continuations of the application that ultimately issued as the '8-296 Patent.

105. ISA-L comprises a collection of optimized low-level functions used for storage applications.

106. ISA-L is optimized for Intel architecture Intel® 64.

107. ISA-L is packaged and shipped with Cloudera CDH.

108. Intel collaborated with Cloudera to apply erasure coding, including on changes made to the NameNode, DataNode, and the client read and write paths, as well as optimizations using Intel ISA-L to accelerate the encoding and decoding calculations.

109. ISA-L is enabled by default in Cloudera CDH.

110. Intel is doing business in the United States and more particularly in this District by actively inducing infringement by at least Cloudera, ADP, Experian, and Wargaming of StreamScale's Patents-in-Suit through the deployment and support of ISA-L.

### **COUNT 1—INFRINGEMENT OF THE '8-296 PATENT**

111. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–110 of this Complaint as though fully set forth herein.

#### **I. DIRECT INFRINGEMENT**

112. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '8-296 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.



113. The EC System Defendants are infringing claims of the '8-296 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

114. Claim 1 of the '8-296 Patent is directed to a system for accelerated error-correcting code (ECC) processing comprising: a processing core for executing computer instructions and accessing data from a main memory; and a non-volatile storage medium for storing the computer instructions, wherein the processing core, the non-volatile storage medium, and the computer instructions are configured to implement an erasure coding system comprising: a data matrix for holding original data in the main memory; a check matrix for holding check data in the main memory; an encoding matrix for holding first factors in the main memory, the first factors being for encoding the original data into the check data; and a thread for executing on the processing core and comprising: a parallel multiplier for concurrently multiplying multiple data entries of a matrix by a single factor; and a first sequencer for ordering operations through the data matrix and the encoding matrix using the parallel multiplier to generate the check data.

**A. CLOUDERA'S DIRECT INFRINGEMENT**

115. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '8-296 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Cloudera Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Cloudera Infringing Products and Services to encode

the original data into check data. The Cloudera Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**B. ADP'S DIRECT INFRINGEMENT**

116. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '8-296 Patent, including at least Claim 1. The ADP Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The ADP Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the ADP Infringing Products and Services to encode the original data into check data. The ADP Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**C. EXPERIAN'S DIRECT INFRINGEMENT**

117. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '8-296 Patent, including at least Claim 1. The Experian Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Experian Infringing Products and Services include non-volatile storage (memory) and computer

instructions to implement accelerated ECC. The accelerated ECC system of the Experian Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Experian Infringing Products and Services to encode the original data into check data. The Experian Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**D. WARGAMING'S DIRECT INFRINGEMENT**

118. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '8-296 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Wargaming Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Wargaming Infringing Products and Services to encode the original data into check data. The Wargaming Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

## II. INDIRECT INFRINGEMENT

119. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '8-296 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '8-296 Patent by third parties, including for example Cloudera, ADP, Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '8-296 Patent.

120. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '8-296 Patent. For example, Intel induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '8-296 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '8-296 Patent. Intel (1) knew or should have known of the '8-296 Patent since at least 2014, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '8-296 Patent's claims by users of ISA-L.

121. Intel actively markets and instructs the EC System Defendants to create EC Systems using ISA-L.

122. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>9</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use that are available to the EC System Defendants on the Intel website,<sup>10</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API Reference Manual for ISA-L<sup>11</sup> that is available to the EC System Defendants, which it updates regularly.<sup>12</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

123. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

124. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '8-296 Patent,

---

<sup>9</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, *available at* <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

<sup>10</sup> *See, e.g.*, Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), *available at* <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

<sup>11</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), *available at* [https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).

<sup>12</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), *available at* [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

including at least Claim 1. EC Systems that employ ISA-L create a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used to encode the original data into check data. The systems also include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

125. As explained above, Intel had actual notice of the '8-296 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '8-296 Patent by Cloudera, ADP, Experian, and Wargaming.

126. Intel further had actual notice of its infringement, and the role of ISA-L in that infringement, of the '8-296 Patent. On or about July 20, 2015, StreamScale sent a claim chart to Intel detailing how ISA-L could be used to infringe at least Claim 1 of the '8-296 Patent.

127. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's patents, including the '8-296 Patent. To the extent that Intel lacked actual knowledge of the '8-296 Patent or the EC System Defendants' actual infringement of the '8-296 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

128. At a minimum, Intel has had actual notice of the '8-296 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '8-296 Patent by Cloudera, ADP, Experian, and Wargaming.

129. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '8-296 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

130. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

### **COUNT 2—INFRINGEMENT OF THE '374 PATENT**

131. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–130 of this Complaint as though fully set forth herein.

#### **I. DIRECT INFRINGEMENT**

132. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '374 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.

133. The EC System Defendants are infringing claims of the '374 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

134. Claim 1 of the '374 Patent is directed to a system for accelerated error-correcting code (ECC) processing comprising: a processing core for executing computer instructions and accessing data from a main memory, the processing core comprising at least 16 data registers, each of the data registers comprising at least 16 bytes; and a non-volatile storage medium for storing the computer instructions, wherein the processing core, the non-volatile storage medium, and the computer instructions are configured to implement an erasure coding system comprising: a data matrix for holding original data in the main memory; a check matrix for holding check data in the main memory; an encoding matrix for holding first factors in the main memory, the first factors being for encoding the original data into the check data; and a thread for executing on the processing core and comprising: a parallel multiplier for concurrently multiplying multiple data entries of a matrix by a single factor; and a first sequencer for ordering operations through the data matrix and the encoding matrix using the parallel multiplier to generate the check data.

**A. CLOUDERA'S DIRECT INFRINGEMENT**

135. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '374 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Cloudera Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an



encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Cloudera Infringing Products and Services to encode the original data into check data. The Cloudera Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**B. ADP'S DIRECT INFRINGEMENT**

136. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '374 Patent, including at least Claim 1. The ADP Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The ADP Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the ADP Infringing Products and Services to encode the original data into check data. The ADP Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**C. EXPERIAN'S DIRECT INFRINGEMENT**

137. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '374 Patent, including at least Claim 1. The Experian Infringing Products

and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Experian Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the Experian Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Experian Infringing Products and Services to encode the original data into check data. The Experian Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

**D. WARGAMING'S DIRECT INFRINGEMENT**

138. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '374 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Wargaming Infringing Products and Services include non-volatile storage (memory) and computer instructions to implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Wargaming Infringing Products and Services to encode the original data into check data. The Wargaming Infringing Products and Services include a thread

for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

## **II. INDIRECT INFRINGEMENT**

139. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '374 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '374 Patent by third parties, including for example Cloudera, ADP, Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '374 Patent.

140. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '374 Patent. For example, Intel induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '374 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '374 Patent. Intel (1) knew or should have known of the '374 Patent since at least 2015, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '374 Patent's claims by users of ISA-L.

141. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>13</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use that are available to the EC System Defendants on the Intel website,<sup>14</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API Reference Manual for ISA-L<sup>15</sup> that is available to the EC System Defendants, which it updates regularly.<sup>16</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

142. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

143. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '374 Patent,

---

<sup>13</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, *available at* <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

<sup>14</sup> *See, e.g.*, Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), *available at* <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

<sup>15</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), *available at* [https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).

<sup>16</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), *available at* [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

including at least Claim 1. EC Systems that employ ISA-L create a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used to encode the original data into check data. The systems also include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering operations through the data matrix and encoding matrix to generate the check data.

144. As explained above, Intel had actual notice of the '374 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '374 Patent by Cloudera, ADP, Experian, and Wargaming.

145. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's patents, including the '374 Patent. To the extent that Intel lacked actual knowledge of the '374 Patent or the EC System Defendants' actual infringement of the '374 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

146. At a minimum, Intel has had actual notice of the '374 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '374 Patent by Cloudera, ADP, Experian, and Wargaming.

147. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '374 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased

damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

148. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

#### **COUNT 3—INFRINGEMENT OF THE '759 PATENT**

149. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–148 of this Complaint as though fully set forth herein.

### **I. DIRECT INFRINGEMENT**

150. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '759 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.

151. The EC System Defendants are infringing claims of the '759 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

152. Claim 1 of the '759 Patent is directed to a system for accelerated error-correcting code (ECC) processing comprising: a processing core for executing computer instructions and accessing data from a main memory, the processing core comprising at least 16 data registers, each of the data registers comprising at least 16 bytes; one or more non-volatile storage media for storing the computer instructions and the data; and an input/output (I/O) controller for controlling data

transfers between the main memory and the non-volatile storage media, wherein the processing core, the non-volatile storage media, the I/O controller, and the computer instructions are configured to implement an erasure coding system comprising: a data matrix for holding original data in the main memory; a check matrix for holding check data in the main memory; an encoding matrix for holding first factors in the main memory, the first factors being for encoding the original data into the check data; and a thread for executing on the processing core and comprising: a parallel multiplier for concurrently multiplying multiple data entries of a matrix by a single factor; and a first sequencer for ordering data accesses through the data matrix and the encoding matrix using the parallel multiplier to generate the check data.

**A. CLOUDERA'S DIRECT INFRINGEMENT**

153. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '759 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Cloudera Infringing Products and Services include non-volatile storage (memory) for storing computer instructions and data. The Cloudera Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media. The processing core, memory, I/O controller, and computer instructions of the Cloudera Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Cloudera Infringing Products and

Services to encode the original data into check data. The Cloudera Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering data accesses through the data matrix and encoding matrix to generate the check data.

## **B. ADP'S DIRECT INFRINGEMENT**

154. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '759 Patent, including at least Claim 1. The ADP Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The ADP Infringing Products and Services include non-volatile storage (memory) for storing computer instructions and data. The ADP Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media. The processing core, memory, I/O controller, and computer instructions of the ADP Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the ADP Infringing Products and Services to encode the original data into check data. The ADP Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering data accesses through the data matrix and encoding matrix to generate the check data.



**C. EXPERIAN'S DIRECT INFRINGEMENT**

155. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '759 Patent, including at least Claim 1. The Experian Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Experian Infringing Products and Services include non-volatile storage (memory) for storing computer instructions and data. The Experian Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media. The processing core, memory, I/O controller, and computer instructions of the Experian Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Experian Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Experian Infringing Products and Services to encode the original data into check data. The Experian Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering data accesses through the data matrix and encoding matrix to generate the check data.

**D. WARGAMING'S DIRECT INFRINGEMENT**

156. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '759 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are systems capable of performing accelerated ECC. They comprise a

processing core comprising at least 16 data registers of at least 16 bytes each, including, for example, one or more Intel, AMD, ARM, and/or PPC64 processing cores. The Wargaming Infringing Products and Services include non-volatile storage (memory) for storing computer instructions and data. The Wargaming Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media. The processing core, memory, I/O controller, and computer instructions of the Wargaming Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used in the Wargaming Infringing Products and Services to encode the original data into check data. The Wargaming Infringing Products and Services include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering data accesses through the data matrix and encoding matrix to generate the check data.

## **II. INDIRECT INFRINGEMENT**

157. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '759 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '759 Patent by third parties, including for example Cloudera, ADP, Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '759 Patent.

158. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive

materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '759 Patent. For example, Intel induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '759 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '759 Patent. Intel (1) knew or should have known of the '759 Patent since at least 2016, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '759 Patent's claims by users of ISA-L.

159. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>17</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use that are available to the EC System Defendants on the Intel website,<sup>18</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API Reference Manual for ISA-L<sup>19</sup> that is available to the EC System Defendants, which it updates

---

<sup>17</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, available at <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

<sup>18</sup> *See, e.g.*, Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L, available at <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

<sup>19</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), available at

regularly.<sup>20</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

160. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

161. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '759 Patent, including at least Claim 1. EC Systems that employ ISA-L create a data matrix for holding original data, a check matrix for holding check data, and an encoding matrix for holding first factors, all in memory. The first factors of the encoding matrix are used to encode the original data into check data. The systems also include a thread for executing on the processing core that includes a parallel lookup multiplier and a sequencer for ordering data accesses through the data matrix and encoding matrix to generate the check data.

162. As explained above, Intel had actual notice of the '759 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '759 Patent by Cloudera, ADP, Experian, and Wargaming.

163. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's

---

[https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).

<sup>20</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), available at [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

patents, including the '759 Patent. To the extent that Intel lacked actual knowledge of the '759 Patent or the EC System Defendants' actual infringement of the '759 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

164. At a minimum, Intel has had actual notice of the '759 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '759 Patent by Cloudera, ADP, Experian, and Wargaming.

165. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '759 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

166. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

#### **COUNT 4—INFRINGEMENT OF THE '358 PATENT**

167. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–166 of this Complaint as though fully set forth herein.

### **I. DIRECT INFRINGEMENT**

168. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '358 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing

into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.

169. The EC System Defendants are infringing claims of the '358 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

170. Claim 1 of the '358 Patent is directed to a system adapted to use accelerated error-correcting code (ECC) processing to improve the storage and retrieval of digital data distributed across a plurality of drives, comprising: at least one processor comprising at least one single-instruction-multiple-data (SIMD) central processing unit (CPU) core that executes SIMD instructions and loads original data from a main memory and stores check data to the main memory, the SIMD CPU core comprising at least 16 vector registers, each of the vector registers storing at least 16 bytes; at least one system drive comprising at least one non-volatile storage medium that stores the SIMD instructions; a plurality of data drives each comprising at least one non-volatile storage medium that stores at least one block of the original data, the at least one block comprising at least 512 bytes; more than two check drives each comprising at least one non-volatile storage medium that stores at least one block of the check data; and at least one input/output (I/O) controller that stores the at least one block of the check data from the main memory to the check drives, wherein the processor, the SIMD instructions, the non-volatile storage media, and the I/O controller are configured to implement an erasure coding system comprising: a data matrix comprising at least one vector and comprising a plurality of rows of at least one block of the original data in the main memory, each of the rows being stored on a different one of the data drives; a check matrix comprising more than two rows of the at least one block of the check data

in the main memory, each of the rows being stored on a different one of the check drives, one of the rows comprising a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix; a thread that executes on the SIMD CPU core and comprising: at least one parallel multiplier that multiplies the at least one vector of the data matrix by a single factor to compute parallel multiplier results comprising at least one vector; at least one parallel adder that adds the at least one vector of the parallel multiplier results and computes a running total; and a sequencer wherein the sequencer orders load operations of the original data into at least one of the vector registers and computes the check data with the parallel lookup multiplier and the parallel adder, and stores the computed check data from the vector registers to the main memory.

**A. CLOUDERA'S DIRECT INFRINGEMENT**

171. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '358 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Cloudera Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Cloudera Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Cloudera Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Cloudera Infringing Products and

Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media and that stores the check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Cloudera Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the Cloudera Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Cloudera Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Cloudera Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

## **B. ADP'S DIRECT INFRINGEMENT**

172. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '358 Patent, including at least Claim 1. The ADP Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The



SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The ADP Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The ADP Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The ADP Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The ADP Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media and that stores the check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the ADP Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the ADP Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The ADP Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the ADP Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

**C. EXPERIAN'S DIRECT INFRINGEMENT**

173. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '358 Patent, including at least Claim 1. The Experian Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Experian Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Experian Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Experian Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Experian Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile storage media and that stores the check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Experian Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Experian Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the Experian Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity

row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Experian Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Experian Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

#### **D. WARGAMING'S DIRECT INFRINGEMENT**

174. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '358 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Wargaming Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Wargaming Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Wargaming Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Wargaming Infringing Products and Services further include an input/output (I/O) controller to coordinate communication and data transfers between the main memory and the non-volatile

storage media and that stores the check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Wargaming Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the Wargaming Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Wargaming Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Wargaming Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

## **II. INDIRECT INFRINGEMENT**

175. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '358 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '358 Patent by third parties, including for example Cloudera, ADP, Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '358 Patent.

176. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '358 Patent. For example, Intel induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '358 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '358 Patent. Intel (1) knew or should have known of the '358 Patent since at least 2018, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '358 Patent's claims by users of ISA-L.

177. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>21</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use that are available to the EC System Defendants on the Intel website,<sup>22</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API

---

<sup>21</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, *available at* <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

<sup>22</sup> *See, e.g.*, Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L, *available at* <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

Reference Manual for ISA-L<sup>23</sup> that is available to the EC System Defendants, which it updates regularly.<sup>24</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

178. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

179. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '358 Patent, including at least Claim 1. EC Systems that employ ISA-L create a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The systems that employ ISA-L create a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The systems also include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The systems' parallel lookup multiplier multiplies a vector of the data matrix by a single factor; the systems' parallel adder adds the result of the parallel multiplier to compute a running total; and the

---

<sup>23</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), available at [https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).

<sup>24</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), available at [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

systems' sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

180. As explained above, Intel had actual notice of the '358 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '358 Patent by Cloudera, ADP, Experian, and Wargaming.

181. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's patents, including the '358 Patent. To the extent that Intel lacked actual knowledge of the '358 Patent or the EC System Defendants' actual infringement of the '358 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

182. At a minimum, Intel has had actual notice of the '358 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '358 Patent by Cloudera, ADP, Experian, and Wargaming.

183. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '358 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

184. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

#### **COUNT 5—INFRINGEMENT OF THE '259 PATENT**

185. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–184 of this Complaint as though fully set forth herein.

#### **I. DIRECT INFRINGEMENT**

186. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '259 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.

187. The EC System Defendants are infringing claims of the '259 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

188. Claim 1 of the '259 Patent is directed to a system adapted to use accelerated error-correcting code (ECC) processing to improve the storage and retrieval of digital data distributed across a plurality of drives, comprising: at least one processor comprising at least one single-instruction-multiple-data (SIMD) central processing unit (CPU) core that executes SIMD instructions and loads original data from a main memory and stores check data to the main memory, the SIMD CPU core comprising at least 16 vector registers, each of the vector registers storing at least 16 bytes; at least one system drive comprising at least one non-volatile storage



medium that stores the SIMD instructions; a plurality of data drives each comprising at least one non-volatile storage medium that stores at least one block of the original data, the at least one block comprising at least 512 bytes; more than two check drives each comprising at least one non-volatile storage medium that stores at least one block of the check data; at least one first input/output (I/O) controller that receives the at least one block of the original data from a transmitter and that stores the at least one block of the original data to the main memory; and at least one second input/output (I/O) controller that stores the at least one block of the check data from the main memory to the check drives, wherein the processor, the SIMD instructions, the non-volatile storage medium, and the at least one second I/O controller are configured to implement an erasure coding system comprising: a data matrix comprising at least one vector and comprising a plurality of rows of at least one block of the original data in the main memory, each of the rows being stored on a different one of the data drives; a check matrix comprising more than two rows of the at least one block of the check data in the main memory, each of the rows being stored on a different one of the check drives, one of the rows comprising a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix; and a thread that executes on the SIMD CPU core and comprising: at least one parallel multiplier that multiplies the at least one vector of the data matrix by a single factor to compute parallel multiplier results comprising at least one vector; at least one parallel adder that adds the at least one vector of the parallel multiplier results and computes a running total; and a sequencer wherein the sequencer orders load operations of the original data into at least one of the vector registers and computes the check data with the parallel multiplier and the parallel adder, and stores the computed check data from the vector registers to the main memory.

**A. CLOUDERA'S DIRECT INFRINGEMENT**

189. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '259 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Cloudera Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Cloudera Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Cloudera Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Cloudera Infringing Products and Services further include a first input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The Cloudera Infringing Products and Services further include a second I/O controller to store blocks of check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Cloudera Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the Cloudera Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different

check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Cloudera Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Cloudera Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

## **B. ADP'S DIRECT INFRINGEMENT**

190. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '259 Patent, including at least Claim 1. The ADP Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The ADP Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The ADP Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The ADP Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The ADP Infringing Products and Services further include a first input/output (I/O) controller to receive blocks of original data from a transmitter and

store that data to the main memory. The ADP Infringing Products and Services further include a second I/O controller to store blocks of check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the ADP Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the ADP Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The ADP Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the ADP Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

### **C. EXPERIAN'S DIRECT INFRINGEMENT**

191. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '259 Patent, including at least Claim 1. The Experian Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory.

The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Experian Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Experian Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Experian Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Experian Infringing Products and Services further include a first input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The Experian Infringing Products and Services further include a second I/O controller to store blocks of check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Experian Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Experian Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The accelerated ECC system of the Experian Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Experian Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Experian Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer

orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

**D. WARGAMING'S DIRECT INFRINGEMENT**

192. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '259 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are systems adapted to use accelerated ECC processing to improve the storage and retrieval of digital data that is distributed across multiple drives. They comprise a processing core comprising a single-instruction-multiple-data ("SIMD") central processing unit ("CPU") core that executes the SIMD instructions and loads data from main memory and stores data to main memory. The SIMD CPU core, including for example Intel, AMD, ARM, and/or PPC64 processing cores, includes at least 16 data registers of at least 16 bytes each. The Wargaming Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the SIMD computer instructions. The Wargaming Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data that are at least 512 bytes. The Wargaming Infringing Products and Services include more than two check drives, each of which includes a memory that stores blocks of check data. The Wargaming Infringing Products and Services further include a first input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The Wargaming Infringing Products and Services further include a second I/O controller to store blocks of check data from the main memory to the check drives. The processing core, SIMD instructions, memory, and I/O controller of the Wargaming Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes a data matrix for holding vectors of original data, with each row of

a block of original data stored on a different data drive. The accelerated ECC system of the Wargaming Infringing Products and Services includes a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The Wargaming Infringing Products and Services include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The parallel lookup multiplier of the Wargaming Infringing Products and Services multiplies a vector of the data matrix by a single factor; the parallel adder adds the result of the parallel multiplier to compute a running total; and the sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

## **II. INDIRECT INFRINGEMENT**

193. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '259 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '259 Patent by third parties, including for example Cloudera, ADP, Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '259 Patent.

194. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '259 Patent. For example, Intel induced

Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '259 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '259 Patent. Intel (1) knew or should have known of the '259 Patent since at least 2019, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '259 Patent's claims by users of ISA-L.

195. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>25</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use that are available to the EC System Defendants on the Intel website,<sup>26</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API Reference Manual for ISA-L<sup>27</sup> that is available to the EC System Defendants, which it updates

---

<sup>25</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, *available at* <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

<sup>26</sup> *See, e.g.*, Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), *available at* <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

<sup>27</sup> *See, e.g.*, Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), *available at* [https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).



regularly.<sup>28</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

196. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

197. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '259 Patent, including at least Claim 1. EC Systems that employ ISA-L create a data matrix for holding vectors of original data, with each row of a block of original data stored on a different data drive. The systems that employ ISA-L create a check matrix for holding vectors of check data, with each row of a block of check data stored on different check drives. Moreover, one of the rows of the block of check data comprises a parity row comprising the Galois Field (GF) summation of all of the rows of the data matrix. The systems also include a thread for executing on the SIMD CPU processing core that includes a parallel lookup multiplier, a parallel adder, and a sequencer. The systems' parallel lookup multiplier multiplies a vector of the data matrix by a single factor; the systems' parallel adder adds the result of the parallel multiplier to compute a running total; and the systems' sequencer orders load operations of the data into the registers, computes the check data, and stores the computed check data to main memory.

198. As explained above, Intel had actual notice of the '259 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '259 Patent by Cloudera, ADP, Experian, and Wargaming.

---

<sup>28</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), available at [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

199. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's patents, including the '259 Patent. To the extent that Intel lacked actual knowledge of the '259 Patent or the EC System Defendants' actual infringement of the '259 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

200. At a minimum, Intel has had actual notice of the '259 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '259 Patent by Cloudera, ADP, Experian, and Wargaming.

201. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '259 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate, consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

202. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

#### **COUNT 6—INFRINGEMENT OF THE '10-296 PATENT**

203. StreamScale incorporates by reference the allegations set forth in Paragraphs 1–202 of this Complaint as though fully set forth herein.

## **I. DIRECT INFRINGEMENT**

204. In violation of 35 U.S.C. § 271(a), Cloudera, ADP, Experian, and Wargaming are and have been directly infringing one or more of the '10-296 Patent's claims, including at least Claim 1, by making, using, selling, and/or offering for sale in the United States, and/or importing into the United States, without authority, erasure code products and services, including but not limited to those utilizing ISA-L, including without limitation the Cloudera Infringing Products and Services, the ADP Infringing Products and Services, the Experian Infringing Products and Services, and the Wargaming Infringing Products and Services, as described above.

205. The EC System Defendants are infringing claims of the '10-296 Patent, including at least Claim 1, literally and/or pursuant to the doctrine of equivalents.

206. Claim 1 of the '10-296 Patent is directed to an accelerated error-correcting code (ECC) system operating across multiple drives, comprising: at least one processing circuit comprising a plurality of central processing unit (CPU) cores that executes CPU instructions and loads original data from a main memory and stores check data to the main memory, each of the CPU cores comprising at least 16 registers, and each of the registers storing at least 8 bytes; at least one system drive comprising at least one non-volatile storage medium that stores the CPU instructions; a plurality of data drives each comprising at least one non-volatile storage medium that stores at least one block of the original data; at least four check drives each comprising at least one non-volatile storage medium that stores at least one block of the check data corresponding to the at least one block of the original data; and at least one input/output (I/O) controller that receives the at least one block of the original data from a transmitter and that stores the at least one block of the original data to a main memory; wherein the processing circuit, the CPU instructions, the main memory, the plurality of data drives, the at least four check drives, and the at least one I/O controller are configured to implement a multi-core erasure encoding system comprising: original

data in the main memory comprised of the at least one block of original data from the plurality of data drives; check data in the main memory comprised of the at least one block of check data; an encoding matrix for holding first factors in the main memory, the first factors being for encoding the original data in the main memory into the check data in the main memory; and a scheduler for generating ECC data in parallel across a plurality of threads by: dividing the original data in the main memory into a plurality of data matrices; dividing the check data in the main memory into a plurality of check matrices; assigning corresponding ones of the data matrices and the check matrices in the main memory to the plurality of threads, wherein each thread comprises an encoder, the encoder comprising at least a portion of the encoding matrix, a Galois Field (GF) multiplier, a Galois Field (GF) adder, and a sequencer for ordering operations through at least one of the data matrices, corresponding ones of the check matrices, and the at least a portion of the encoding matrix in the main memory using the GF multiplier and the GF adder to generate the check data in the main memory; and assigning the plurality of threads to the plurality of CPU cores of the processing circuit to concurrently generate the check matrices in the main memory from corresponding ones of the data matrices in the main memory.

**A. CLOUDERA’S DIRECT INFRINGEMENT**

207. As to Cloudera, at least the Cloudera Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the ’10-296 Patent, including at least Claim 1. The Cloudera Infringing Products and Services are accelerated ECC systems operating across multiple drives. They comprise a processing circuit comprising multiple central processing unit (“CPU”) cores that execute CPU instructions and loads original data from main memory and stores check data to main memory. The CPU processing cores, including for example Intel, AMD, ARM, and/or PPC64 processing cores, each include at least 16 data registers of at least 8 bytes each. The Cloudera Infringing

Products and Services include a system drive with non-volatile storage (memory) for storing the CPU instructions. The Cloudera Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data. The Cloudera Infringing Products and Services include at least four check drives, each of which includes a memory that stores blocks of check data, each block of check data corresponding to a block of the original data. The Cloudera Infringing Products and Services further include an input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The processing circuit, CPU instructions, memory, data drives, check drives, and I/O controller of the Cloudera Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Cloudera Infringing Products and Services includes original data blocks in main memory from the multiple data drives, check data blocks in main memory, and encoding matrix with first factors in main memory where the first factors are for encoding the original data into check data, and a scheduler that generates ECC data in parallel across multiple threads. The scheduler of the Cloudera Infringing Products and Services divides the original data in main memory into multiple data matrices and the check data in main memory into multiple check matrices. The scheduler of the Cloudera Infringing Products and Services further assigns data and check matrices to the threads. Each thread in the Cloudera Infringing Products and Services includes an encoder comprising at least part of the encoding matrix, a Galois Field (GF) multiplier, a GF adder, and a sequencer that orders operations of the data to generate the check data in the main memory. The scheduler of the Cloudera Infringing Products and Services further assigns the threads to the various CPU cores of the processing circuit to concurrently generate the check matrices from the data matrices in main memory.

**B. ADP'S DIRECT INFRINGEMENT**

208. As to ADP, at least the ADP Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '10-296 Patent, including at least Claim 1. The ADP Infringing Products and Services are accelerated ECC systems operating across multiple drives. They comprise a processing circuit comprising multiple central processing unit ("CPU") cores that execute CPU instructions and loads original data from main memory and stores check data to main memory. The CPU processing cores, including for example Intel, AMD, ARM, and/or PPC64 processing cores, each include at least 16 data registers of at least 8 bytes each. The ADP Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the CPU instructions. The ADP Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data. The ADP Infringing Products and Services include at least four check drives, each of which includes a memory that stores blocks of check data, each block of check data corresponding to a block of the original data. The ADP Infringing Products and Services further include an input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The processing circuit, CPU instructions, memory, data drives, check drives, and I/O controller of the ADP Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the ADP Infringing Products and Services includes original data blocks in main memory from the multiple data drives, check data blocks in main memory, and encoding matrix with first factors in main memory where the first factors are for encoding the original data into check data, and a scheduler that generates ECC data in parallel across multiple threads. The scheduler of the ADP Infringing Products and Services divides the original data in main memory into multiple data matrices and the check data in main memory into multiple check matrices. The scheduler of the ADP Infringing

Products and Services further assigns data and check matrices to the threads. Each thread in the ADP Infringing Products and Services includes an encoder comprising at least part of the encoding matrix, a Galois Field (GF) multiplier, a GF adder, and a sequencer that orders operations of the data to generate the check data in the main memory. The scheduler of the ADP Infringing Products and Services further assigns the threads to the various CPU cores of the processing circuit to concurrently generate the check matrices from the data matrices in main memory.

### **C. EXPERIAN'S DIRECT INFRINGEMENT**

209. As to Experian, at least the Experian Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '10-296 Patent, including at least Claim 1. The Experian Infringing Products and Services are accelerated ECC systems operating across multiple drives. They comprise a processing circuit comprising multiple central processing unit ("CPU") cores that execute CPU instructions and loads original data from main memory and stores check data to main memory. The CPU processing cores, including for example Intel, AMD, ARM, and/or PPC64 processing cores, each include at least 16 data registers of at least 8 bytes each. The Experian Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the CPU instructions. The Experian Infringing Products and Services include multiple data drives, each of which includes a memory that stores blocks of original data. The Experian Infringing Products and Services include at least four check drives, each of which includes a memory that stores blocks of check data, each block of check data corresponding to a block of the original data. The Experian Infringing Products and Services further include an input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The processing circuit, CPU instructions, memory, data drives, check drives, and I/O controller of the Experian Infringing Products and Services implement accelerated ECC. The accelerated ECC

system of the Experian Infringing Products and Services includes original data blocks in main memory from the multiple data drives, check data blocks in main memory, and encoding matrix with first factors in main memory where the first factors are for encoding the original data into check data, and a scheduler that generates ECC data in parallel across multiple threads. The scheduler of the Experian Infringing Products and Services divides the original data in main memory into multiple data matrices and the check data in main memory into multiple check matrices. The scheduler of the Experian Infringing Products and Services further assigns data and check matrices to the threads. Each thread in the Experian Infringing Products and Services includes an encoder comprising at least part of the encoding matrix, a Galois Field (GF) multiplier, a GF adder, and a sequencer that orders operations of the data to generate the check data in the main memory. The scheduler of the Experian Infringing Products and Services further assigns the threads to the various CPU cores of the processing circuit to concurrently generate the check matrices from the data matrices in main memory.

**D. WARGAMING'S DIRECT INFRINGEMENT**

210. As to Wargaming, at least the Wargaming Infringing Products and Services, as defined above, comprise hardware and software components that together practice every element of one or more claims of the '10-296 Patent, including at least Claim 1. The Wargaming Infringing Products and Services are accelerated ECC systems operating across multiple drives. They comprise a processing circuit comprising multiple central processing unit ("CPU") cores that execute CPU instructions and loads original data from main memory and stores check data to main memory. The CPU processing cores, including for example Intel, AMD, ARM, and/or PPC64 processing cores, each include at least 16 data registers of at least 8 bytes each. The Wargaming Infringing Products and Services include a system drive with non-volatile storage (memory) for storing the CPU instructions. The Wargaming Infringing Products and Services include multiple



data drives, each of which includes a memory that stores blocks of original data. The Wargaming Infringing Products and Services include at least four check drives, each of which includes a memory that stores blocks of check data, each block of check data corresponding to a block of the original data. The Wargaming Infringing Products and Services further include an input/output (I/O) controller to receive blocks of original data from a transmitter and store that data to the main memory. The processing circuit, CPU instructions, memory, data drives, check drives, and I/O controller of the Wargaming Infringing Products and Services implement accelerated ECC. The accelerated ECC system of the Wargaming Infringing Products and Services includes original data blocks in main memory from the multiple data drives, check data blocks in main memory, and encoding matrix with first factors in main memory where the first factors are for encoding the original data into check data, and a scheduler that generates ECC data in parallel across multiple threads. The scheduler of the Wargaming Infringing Products and Services divides the original data in main memory into multiple data matrices and the check data in main memory into multiple check matrices. The scheduler of the Wargaming Infringing Products and Services further assigns data and check matrices to the threads. Each thread in the Wargaming Infringing Products and Services includes an encoder comprising at least part of the encoding matrix, a Galois Field (GF) multiplier, a GF adder, and a sequencer that orders operations of the data to generate the check data in the main memory. The scheduler of the Wargaming Infringing Products and Services further assigns the threads to the various CPU cores of the processing circuit to concurrently generate the check matrices from the data matrices in main memory.

## **II. INDIRECT INFRINGEMENT**

211. In violation of 35 U.S.C. §§ 271(b), Intel is and has been infringing one or more of the '10-296 Patent's claims, including at least Claim 1, indirectly by inducing the infringement of at least Claim 1 of the '10-296 Patent by third parties, including for example Cloudera, ADP,

Experian, and Wargaming, in this District and elsewhere in the United States. Direct infringement is the result of activities performed by users of systems that incorporate, among other features, ISA-L, including for example Cloudera, ADP, Experian, and Wargaming, in accordance with at least Claim 1 of the '10-296 Patent.

212. Intel's affirmative acts of selling and/or distributing ISA-L (or portions thereof), causing ISA-L (or portions thereof) to be manufactured and distributed, providing instructive materials and information concerning operation and use of ISA-L (or portions thereof), and providing maintenance/service for such products or services, induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '10-296 Patent. For example, Intel induced Cloudera, ADP, Experian, and Wargaming to infringe at least Claim 1 of the '10-296 Patent through the implementation of ISA-L in the Cloudera, ADP, Experian, and Wargaming Infringing Products and Services. By and through these acts, Intel knowingly and specifically intended the users of ISA-L (or portions thereof) to infringe at least Claim 1 of the '10-296 Patent. Intel (1) knew or should have known of the '10-296 Patent since at least 2020, (2) performed and continues to perform affirmative acts that constitute induced infringement, and (3) knew or should have known that those acts would induce actual infringement of one or more of the '10-296 Patent's claims by users of ISA-L.

213. For example, upon information and belief, Intel (i) maintains a website to promote ISA-L,<sup>29</sup> including to the EC System Defendants, (ii) produces videos regarding ISA-L and its use

---

<sup>29</sup> *E.g.*, Intel, Intel® Intelligent Storage Acceleration Library, *available at* <https://software.intel.com/content/www/us/en/develop/tools/isa-l.html> (last visited May 24, 2021).

that are available to the EC System Defendants on the Intel website,<sup>30</sup> (iii) describes case studies on big data optimization using ISA-L that are available to the EC System Defendants on the Intel website, (iv) hosts articles, blog posts, and webinars regarding the use of ISA-L that are available to the EC System Defendants on the Intel website, and (v) publishes and makes available an API Reference Manual for ISA-L<sup>31</sup> that is available to the EC System Defendants, which it updates regularly.<sup>32</sup> Upon information and belief, Intel further offers the EC System Defendants technical support for ISA-L and the EC System Defendants' products.

214. Upon information and belief, Intel promotes and encourages the EC System Defendants to use ISA-L in order to drive sales of other Intel products and services to the EC System Defendants.

215. As to Intel, at least ISA-L, as defined above, is designed to be used with other components that, when combined with hardware, practice one or more claims of the '10-296 Patent, including at least Claim 1. EC Systems that employ ISA-L create original data blocks in main memory from the multiple data drives, check data blocks in main memory, and encoding matrix with first factors in main memory where the first factors are for encoding the original data into check data, and a scheduler that generates ECC data in parallel across multiple

---

<sup>30</sup> See, e.g., Intel, Erasure Code and Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), available at <https://www.intel.com/content/www/us/en/products/docs/storage/erasure-code-isa-l-solution-video.html> (last visited May 24, 2021).

<sup>31</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version, API Reference Manual (ver. 2.8, Sept. 27, 2013), available at [https://01.org/sites/default/files/documentation/isa-l\\_open\\_src\\_2.8\\_0.pdf](https://01.org/sites/default/files/documentation/isa-l_open_src_2.8_0.pdf) (last visited May 24, 2021).

<sup>32</sup> See, e.g., Intel, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L), API Reference Manual (ver. 2.23.0, June 29, 2018), available at [https://01.org/sites/default/files/documentation/isa-l\\_api\\_2.23.0.pdf](https://01.org/sites/default/files/documentation/isa-l_api_2.23.0.pdf) (last visited May 24, 2021).

threads. The scheduler of the systems divides the original and check data into multiple data and check matrices, respectively, and assigns data and check matrices to the threads. Each thread in the systems includes an encoder comprising at least part of the encoding matrix, a Galois Field (GF) multiplier, a GF adder, and a sequencer that orders operations of the data to generate the check data in the main memory.

216. As explained above, Intel had actual notice of the '10-296 Patent prior to this lawsuit being filed and had knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '10-296 Patent by Cloudera, ADP, Experian, and Wargaming.

217. Especially in light of its actual knowledge of StreamScale and StreamScale's patent portfolio, Intel subjectively believed there was a high probability that StreamScale's patents implicated ISA-L and that EC System Defendants use of ISA-L would infringe StreamScale's patents, including the '10-296 Patent. To the extent that Intel lacked actual knowledge of the '10-296 Patent or the EC System Defendants' actual infringement of the '10-296 Patent, Intel took deliberate actions to avoid learning of those facts. Indeed, Intel actively encouraged others to ignore StreamScale and its patents and further reprimanded at least one employee for failing to ignore StreamScale and its patents.

218. At a minimum, Intel has had actual notice of the '10-296 Patent since March 5, 2021 and has knowledge of the infringing nature of its activities, yet continues to induce infringement of at least Claim 1 of the '10-296 Patent by Cloudera, ADP, Experian, and Wargaming.

219. Therefore, upon information and belief, Intel's infringement of at least Claim 1 of the '10-296 Patent has been and continues to be willful, wanton, malicious, bad-faith, deliberate,

consciously wrongful, flagrant, or characteristic of a pirate, entitling StreamScale to increased damages pursuant to 35 U.S.C. § 284 and to attorneys' fees and costs incurred in prosecuting this action pursuant to 35 U.S.C. § 285.

### **III. DAMAGES**

220. Defendants' acts of infringement have caused damages to StreamScale, and StreamScale is entitled to recover from Defendants the damages sustained by StreamScale as a result of Defendants' wrongful acts in an amount to be determined at trial.

#### **DAMAGES**

221. StreamScale is entitled to, and now seeks to, recover damages in an amount not less than the maximum amount permitted by law caused by Defendants' acts of infringement.

222. As a result of Defendants' acts of infringement, StreamScale has suffered actual and consequential damages. To the fullest extent permitted by law, StreamScale seeks recovery of damages in an amount to compensate for Defendants' infringement. StreamScale further seeks any other damages to which StreamScale would be entitled to in law or in equity.

#### **INJUNCTIVE RELIEF**

223. Defendants' acts of infringement have caused—and unless restrained and enjoined, Defendants' acts of infringement will continue to cause—irreparable injury and damage to StreamScale for which StreamScale has no adequate remedy at law. Unless preliminarily and permanently enjoined by this Court, Defendants will continue to infringe the Patents-in-Suit.

#### **ATTORNEYS' FEES**

224. StreamScale is entitled to recover reasonable and necessary attorneys' fees under applicable law.

**DEMAND FOR JURY TRIAL**

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, StreamScale demands a trial by jury on all issues so triable.

**PRAYER FOR RELIEF**

StreamScale respectfully requests that the Court enter preliminary and final orders, declarations, and judgments against Defendants as are necessary to provide StreamScale with the following relief:

- a. A judgment that Defendants have infringed and/or are infringing one or more claims of the '8-296 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- b. A judgment that Intel's infringement of the '8-296 Patent has been willful;
- c. A judgment that Defendants have infringed and/or are infringing one or more claims of the '374 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- d. A judgment that Intel's infringement of the '374 Patent has been willful; A judgment that Defendants have infringed and/or are infringing one or more claims of the '759 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- e. A judgment that Intel's infringement of the '759 Patent has been willful;
- f. A judgment that Defendants have infringed and/or are infringing one or more claims of the '358 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- g. A judgment that Intel's infringement of the '358 Patent has been willful;

- h. A judgment that Defendants have infringed and/or are infringing one or more claims of the '259 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- i. A judgment that Intel's infringement of the '259 Patent has been willful;
- j. A judgment that Defendants have infringed and/or are infringing one or more claims of the '10-296 Patent, literally or under the doctrine of equivalents, and directly or indirectly as alleged above;
- k. A judgment that Intel's infringement of the '10-296 Patent has been willful;
- l. An award for all damages arising out of Defendants' infringement, together with prejudgment and post-judgment interest, jointly and severally, in an amount according to proof, including without limitation attorneys' fees and litigation costs and expenses;
- m. An accounting of damages and any future compensation due to StreamScale for Defendants' infringement (past, present, or future) not specifically accounted for in a damages award (or other relief), and/or permanent injunctive relief;
- n. An award of reasonable attorneys' fees as provided by 35 U.S.C. § 285 and enhanced damages as provided by 35 U.S.C. § 284;
- o. The entry of an order preliminarily and permanently enjoining and restraining Defendants and its parents, affiliates, subsidiaries, officers, agents, servants, employees, attorneys, successors, and assigns and all those person in active concert or participation with them or any of them, from making, importing, using, offering for sale, selling, or causing to be sold

any product falling within the scope of any claim of the Patents-in-Suit, or otherwise infringing or inducing infringement of any claim of the Patents-in-Suit; and

p. All further relief in law or in equity as the Court may deem just and proper.

Dated: May 28, 2021

Respectfully submitted,

/s/ Jamie H. McDole

Jamie H. McDole

State Bar No. 24082049

Phillip B. Philbin

State Bar No. 15909020

Michael D. Karson

State Bar No. 24090198

Nadia E. Haghghatian

State Bar No. 24087652

Austin C. Teng

State Bar No. 24093247

THOMPSON & KNIGHT LLP

One Arts Plaza

1722 Routh St., Suite 1500

Dallas, Texas 75201

Tel.: 214.969.1700

Fax: 214.969.1751

Email: [phillip.philbin@tklaw.com](mailto:phillip.philbin@tklaw.com)

[jamie.mcdole@tklaw.com](mailto:jamie.mcdole@tklaw.com)

[michael.karson@tklaw.com](mailto:michael.karson@tklaw.com)

[nadia.haghghatian@tklaw.com](mailto:nadia.haghghatian@tklaw.com)

[austin.teng@tklaw.com](mailto:austin.teng@tklaw.com)

*Attorneys for Plaintiff StreamScale, Inc.*



**CERTIFICATE OF SERVICE**

I hereby certify that, on May 28, 2021, I electronically submitted the foregoing document with the clerk of the United States District Court for the Western District of Texas, using the electronic case management CM/ECF system of the Court which will send notification of such filing to the following:

Brock S. Weber  
Christopher Kao  
Pillsbury Winthrop Shaw Pittman LLP  
4 Embarcadero Center, 22nd Floor  
San Francisco, CA 94111

Christopher S. Ponder  
Harper Batts  
Sheppard, Mullin, Richter & Hampton LLP  
379 Lytton Avenue  
Palo Alto, CA 94301

Steven P. Tepera  
Pillsbury Winthrop Shaw Pittman LLP  
401 Congress Avenue  
Suite 1700  
Austin, TX 78701

Jennifer Klein Ayers  
Sheppard Mullin Richter & Hampton LLP  
2200 Ross Avenue, 24th Floor  
Dallas, TX 75201

*Counsel for Defendants Cloudera, Inc. &  
Experian plc*

*Counsel for Defendant Wargaming (Austin),  
Inc.*

Amanda L. Major  
Wilmer Cutler Pickering Hale and Dorr LLP  
1875 Pennsylvania Ave., Nw  
Washington, DC 20006

Jose Carlos Villarreal  
Perkins Coie LLP  
500 W. 2nd Street  
Suite 1900  
Austin, TX 78701

Annaleigh E. Curtis  
Wilmer Cutler Pickering Hale and Dorr LLP  
60 State Street  
Boston, MA 02109

*Counsel for Defendant Automatic Data  
Processing, Inc.*

Jennifer J. John  
Sonal N. Mehta  
Wilmer Cutler Pickering Hale and Dorr LLP  
2600 El Camino Real, Suite 400  
Palo Alto, CA 94306

Joseph Taylor Gooch  
Wilmer Cutler Pickering Hale and Dorr LLP  
One Front Street, Suite 3500  
San Francisco, CA 94111

Vikram Iyer  
Wilmer Cutler Pickering Hale and Dorr LLP  
350 South Grand Avenue, Suite 2400  
Los Angeles, CA 90071

Austin Michael Schnell  
Brian Christopher Nash  
Pillsbury Winthrop Shaw Pittman LLP  
401 Congress Ave, Suite 1700  
Austin, TX 78701

*Counsel for Defendant Intel Corp.*

/s/ Jamie H. McDole

Jamie H. McDole