IN THE UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

K.MIZRA LLC, Plaintiff, v. AMAZON.COM, INC. and AMAZON WEB SERVICES, INC., Defendants.	Case No. 1:26-cv-00316 **JURY TRIAL DEMANDED**

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff K.Mizra LLC ("K.Mizra") files this Complaint for patent infringement against Defendants Amazon.com, Inc. ("Amazon.com") and Amazon Web Services, Inc. ("AWS") (collectively, "Defendants" or "Amazon"), alleging as follows:

I. INTRODUCTION

1. K.Mizra is a patent licensing company run by experienced management. The company focuses on high value, high quality patents with a global reach. It owns patent portfolios originating with a wide array of inventors, including portfolios developed by well-known multinationals such as IBM, Intel, Rambus and others, as well as from research institutes such as Nederlandse Organisatie voor Toegespast Natuurwetenschappelijk Onderzoek (Netherlands Organization for Applied Scientific Research, or "TNO"). By focusing on high quality patents, K.Mizra provides a secondary market for inventors to recoup their research and development investments and to continue their innovations. K.Mizra offers licenses to its patents on reasonable terms and in this way plays an important part in the development of the technologies that improve all our lives.

2. K.Mizra remains ready, willing, and able to provide commercially-reasonable licenses for its various patented technologies to all entities who wish or need to use them internally or in connection with products or services offered to others. As outlined below, Amazon is one such entity.

II. THE PARTIES

3. K.Mizra is a Delaware limited liability company with a mailing address of 777 Brickell Avenue, #500-96031, Miami, Florida 33131, and operates in Florida. K.Mizra is the owner by assignment of the Asserted Patents.

4. Defendant Amazon.com, Inc. is a Delaware corporation with a principal place of business at 410 Terry Ave. North, Seattle, Washington 98109-5210. Amazon may be served with process via its registered agent Corporation Service Company, MC-CSC1, 300 Deschutes Way SW, Suite 208, Tumwater, Washington 98501 or via its listed registered agent of Corporation Service Company, 251 Little Falls Drive, Wilmington, Delaware 19808.

5. Defendant Amazon Web Services, Inc. is a Delaware corporation with a principal place of business at 410 Terry Ave. North, Seattle, Washington 98109. Amazon Web Services, Inc. may be served through its registered agent Corporation Service Company, 211 E. 7th Street, Suite 620, Austin, Texas 78701. On information and belief, Amazon Web Services, Inc. is registered to do business in the State of Texas and has been since at least May 3, 2006. On information and belief, Amazon Web Services, Inc. is a wholly-owned subsidiary of Amazon.com, Inc.

III. JURISDICTION AND VENUE

6. This is an action for patent infringement under the patent laws of the United States, 35 U.S.C. §§ 1 *et seq.*, including 35 U.S.C. §§ 271, 281, and 284, among others. The Court has

subject-matter jurisdiction over the claims raised in this action pursuant to 28 U.S.C. §§ 1331 and 1338(a).

7. This Court has personal jurisdiction over Defendants as they have committed acts within this District giving rise to this action and have established minimum contacts with this forum such that the exercise of jurisdiction over Defendants would not offend traditional notions of fair play and substantial justice. Defendants, directly and through subsidiaries or intermediaries, have committed and continue to commit acts of infringement in this District and elsewhere by, among other things, importing, offering to sell, and selling products and/or services that infringe the Asserted Patents.

8. Venue is proper in this District pursuant to 28 U.S.C. § 1400(b) because Defendants have committed acts of infringement in this District and Defendants have a regular and established place of business in this District at Amazon Tech Hub located at 11501 Alterra Parkway, Austin, Texas.

IV. GENERAL ALLEGATIONS

A. Asserted Patents

9. U.S. Patent No. 8,234,705 ("the '705 Patent"), titled "Contagion Isolation and Inoculation," was legally issued by the USPTO to inventors Dr. Roskind and Mr. Emigh on July 31, 2012. A true and correct copy of the '705 Patent is attached hereto as Exhibit 1. The '705 Patent claims priority to U.S. Provisional Application No. 60/613,909, filed on September 27, 2004 ("the '909 Application").

10. U.S. Patent No. 8,144,717 ("the '717 Patent"), titled "Initialization of a Wireless Communication Network," was legally issued by the USPTO to inventors Marinus Johannes Blange and Miodrag Djurica on March 27, 2012. A true and correct copy of the '717 Patent is

attached hereto as Exhibit 2. The '717 Patent claims priority to PCT Application No. PCT/NL2007/000001, filed on January 2, 2007 ("the '001 Application"), and to U.S. Patent Application No. 12/159,753 ("the '753 Application").

11. U.S. Patent No. 8,438,120 ("the '120 Patent"), titled "Machine Learning Hyperparameter Estimation," was legally issued by the USPTO to inventor Stephan Alexander Raaijmakers on May 7, 2013. A true and correct copy of the '120 Patent is attached hereto as Exhibit 3. The '120 Patent claims priority to PCT Application No. PCT/NL2008/050247, filed on April 25, 2008 ("the '247 Application"), and to U.S. Patent Application No. 12/597,257 ("the '257 Application").

12. U.S. Patent No. 9,235,259 ("the '259 Patent"), titled "Method for Detecting Audio Ticks in a Noisy Environment," was legally issued by the USPTO to inventors Mark van Staalduinen, Victor Bastiaan Klos and Peter Jan Otto Doets on January 12, 2016. A true and correct copy of the '259 Patent is attached hereto as Exhibit 4. The '259 Patent claims priority to PCT Application No. PCT/NL2010/050795, filed on November 26, 2010 ("the '795 Application"), and to U.S. Patent Application No. 13/512,139 ("the '139 Application").

13. U.S. Patent No. 9,602,649 ("the '649 Patent"), titled "Event Disambiguation," was legally issued by the USPTO to inventors Jeroen Laarakkers, Mattijs Oskar van Deventer and Victor Bastiaan Klos on March 21, 2017. A true and correct copy of the '649 Patent is attached hereto as Exhibit 5. The '649 Patent claims priority to PCT Application No. PCT/NL2010/050478, filed on July 23, 2010 ("the '478 Application"), and to U.S. Patent Application No. 13/386,461 ("the '461 Application").

14. U.S. Patent No. 8,782,282 ("the '282 Patent"), titled "Network Management System," was legally issued by the USPTO to inventors Eileen Zhou, Roger Liu, Vijoy Choyi,

Moshe Itah and John Z. Yu on July 15, 2014. A true and correct copy of the '282 Patent is attached hereto as Exhibit 6. The '282 Patent claims priority to U.S. Patent Application No. 10/742,573, filed on December 19, 2003 ("the '573 Application").

15. K.Mizra is the sole owner by assignment of the '705 Patent, '717 Patent, '120 Patent, '259 Patent, '649 Patent and '282 Patent (collectively, "the Asserted Patents") with the full and exclusive right to bring suit to enforce them. K.Mizra is also entitled to sue to collect damages for all past infringement of the Asserted Patents.

1. <u>The '705 Patent</u>

16. The '705 Patent was involved in an unsuccessful *Inter Partes* Review Proceeding ("IPR") and several now-resolved federal court litigations, and was originally invented by two highly respected and prolific individual inventors, James A. Roskind and Aaron T. Emigh.

17. The '705 Patent was originally owned by Dr. Roskind and Mr. Emigh's company, Radix Labs, LLC. Dr. Roskind and Mr. Emigh were then, and remain today, focused on innovation, conducting new research, developing new technologies, and creating new and innovative computer products.

18. Dr. Roskind, one of the two inventors of the '705 Patent, has bachelor's, master's, and doctorate degrees from the Massachusetts Institute of Technology in both electrical engineering and computer science, and is the named inventor of over 300 U.S. patents. He has worked for Netscape as the Chief Architect and as the Netcenter Security Architect and was a co-founder for Infoseek, a company that was eventually acquired by Disney for $770 million. He was also a key developer of Google's "transport protocol" that provides the tech giant billions of dollars in value every year.

19. Mr. Emigh, the other named inventor of the '705 Patent, graduated from the University of California, Santa Cruz with degrees in linguistics and computer and information sciences, and is the named inventor of over 140 patents. Prior to working with Dr. Roskind, Mr. Emigh worked in various positions developing software, including working as a software manager, architect, and engineer for Unicom and working as a manager for the software development and technical marketing groups for Philips TriMedia. He has founded or co-founded many companies, in addition to Radix Labs, LLC, including CommerceFlow, Inc., which was acquired by eBay for its technology that Mr. Emigh helped develop.

20. After the '705 Patent issued, Dr. Roskind and Mr. Emigh recouped their research and development investment by selling their rights thereto and continued on in their individual technology development pursuits. K.Mizra ultimately acquired the '705 Patent and licensed it to many of the who's-who of the tech world. Some of the accused infringers chose to test the validity of the '705 Patent before settling their lawsuits involving the '705 Patent. For instance, a few accused infringers of the '705 Patent previously sought IPR by the Patent Trial and Appeal Board ("PTAB"). A Final Written Decision ("Decision") in the IPR found that the petitioners had not shown, by a preponderance of the evidence, that the asserted claims were unpatentable. The IPR Decision was appealed to the U.S. Court of Appeals for the Federal Circuit ("CAFC"), resulting in a procedurally focused remand to the PTAB. Prior to the issuance of the mandate that would have sent the IPR back to the United States Patent and Trademark Office ("USPTO") for further consideration, the parties agreed to move to dismiss the appeal.

a. Prior licensing and litigation of the '705 Patent

21. The '705 Patent has been owned by several entities, in addition to Radix Labs, LLC and K.Mizra, with some of those entities issuing to third parties certain rights to the technologies covered thereby.

22. K.Mizra has been involved in a number of actions it had to institute to protect its patent rights, including actions involving the '705 Patent. Most of those actions resulted in the execution of confidential patent license agreements.

b. Network security problems in 2004 solved by the '705 Patent

23. The technology described in the '705 Patent was invented by Dr. Roskind and Mr. Emigh, two colleagues living in the same area who had similar interests in innovating computer-related technologies. In 2003, the inventors decided to create a business—Radix Labs, LLC—which focused on developing intellectual property related to various computer technologies, including computer network security technologies. The inventors focused on conceiving and reducing to practice inventions that they knew were needed (or soon would be needed) in the computer networking industry and then on drafting patent applications to capture and protect their technological innovations. In September 2004, the inventors filed the '909 Application to which the '705 Patent claims priority. The '909 Application described technology that focused on securing a computer network against the threats to which it was exposed when computer endpoints (e.g., laptop computers) were connected to a computer network. The '909 Application, and by natural extension the '705 Patent, also focuses on remedying identified threats and quarantining those threats to mitigate any damage to the secured network.

24. Claims of the '705 Patent are directed to technological solutions that address specific challenges grounded in computer network security. Maintaining the security of computer

systems and networks is a tremendous concern for modern enterprises, since a breach of an internal network can have severe repercussions, including major financial losses, data theft, disclosure of sensitive information, network disruptions, data corruption, etc. The inventors of the '705 Patent understood that while a network security appliance or hardware can be adept at keeping out unwanted external intrusions from the network, the most exploitable vulnerabilities of most networks are the end-user computers that roam throughout various public and private network domains, potentially exposing those computers to infection and then accessing and potentially infecting the entire and presumably secure computer network.

25. For example, the '705 Patent explains that "[l]aptop and wireless computers and other mobile systems pose a threat to elements comprising and/or connected to a network service provider, enterprise, or other protected network to which they reconnect after a period of connection to one or more networks and/or systems that are not part of the service provider, enterprise, or other protected network. By roaming to unknown domains, such as the Internet, and/or connecting to such domains through public, wireless, and/or otherwise less secure access nodes, such mobile systems may become infected by computer viruses, worms, backdoors, and/or countless other threats and/or exploits and/or have unauthorized software installed; have software installed on the mobile system by an operator of the protected network for the protection of the mobile system and/or the protected network removed or altered without authorization and/or have configurations, settings, security data, and/or other data added, removed, and/or changed in authorized ways and/or by unauthorized person[s]." (*See, e.g.*, Ex. 1 at 1:14–31.)

26. The solution to these problems—as specified and claimed in the '705 Patent—was an advanced departure from the conventional network access control solutions then in use and was then, as it remains today, patent eligible, highly valuable, novel, and non-obvious technology.

c. The '705 Patent claims are presumed valid and patent-eligible

27. K.Mizra asserts that at least, and without limitation, Claim 19 of the '705 Patent has been directly infringed, either literally or under the doctrine of equivalents. K.Mizra reserves the right to assert additional claims of the '705 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '705 Claims."

28. None of the Asserted '705 Claims are directed to abstract ideas, and each employs inventive concepts and is directed to patent-eligible subject matter. All claims of the '705 Patent are also presumed to be valid and enforceable against Amazon and others.

29. Indeed, the '705 Patent's specification and claims demonstrate that the need satisfied by the inventions of the Asserted '705 Claims was long-felt in the industry and thus unconventional. As one example, the '705 Patent explains that mobile end user devices such as laptops and wireless computers pose a threat to protected network elements because those devices may access unsecure systems and thereby "become infected by computer viruses, worms, backdoors, and/or countless other threats and/or exploits and/or have unauthorized software installed; have software installed on the mobile system by an operator of the protected network for the protection of the mobile system and/or the protected network removed or altered without authorization; and/or have configurations, settings, security data, and/or other data added, removed, and/or changed in unauthorized ways and/or by unauthorized person[s]." (Ex. 1 at 1:23-31.) Similarly, stationary systems such as desktop computers "may become infected, e.g., due to receipt and execution of malicious code via a network or other communication and/or a diskette and/or other removable media." (*Id.* at 1:31-34.) This poses a danger to a protected network because, when the user device connects to the protected network, "a system may infect or otherwise

harm resources associated with the protected network before measures can be taken to detect and prevent the spread of such infections or harm." (*Id.* at 1:34-38.) The '705 Patent's specification further provides that "[t]herefore, there is a need for a reliable way to ensure that a system does not infect or otherwise harm other network resources when connected to a protected network." (*Id.* at 1:38-41.)

30. The specification (including the provisions quoted above), the figures (including those included below), and the text related to the figures further illustrate the complex, tiered network system architecture of the inventions captured by the Asserted '705 Claims. These figures include the following:

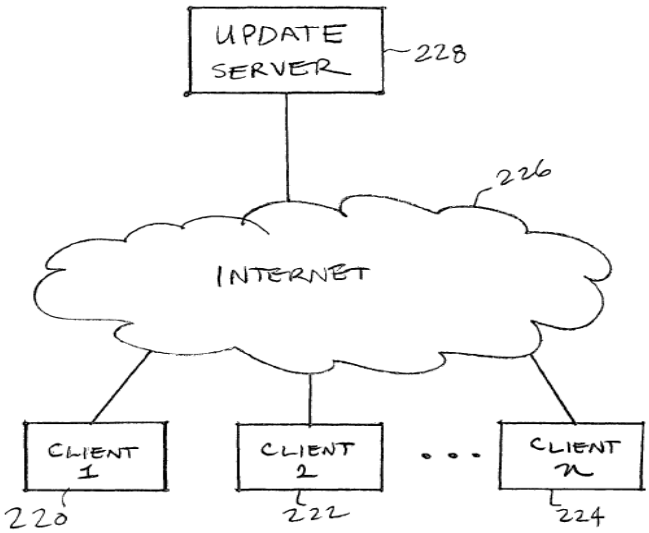

FIG. 2B

(*See id.* at Fig. 2B.)

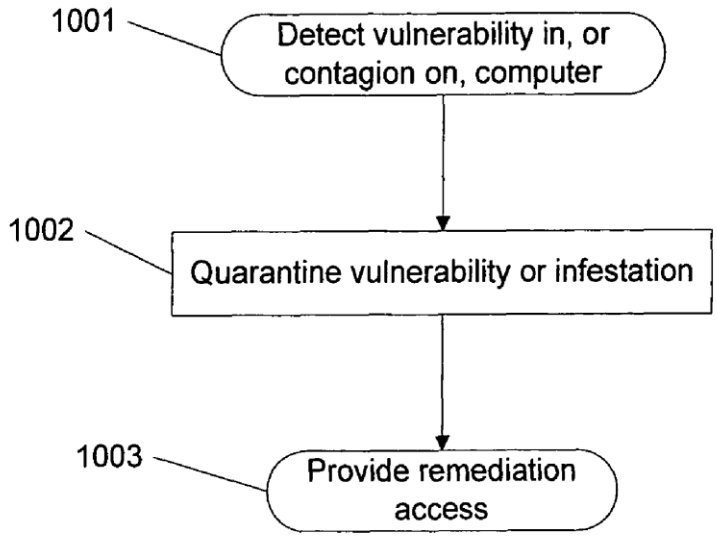

(*See id.* at Fig. 10A.)

31. The foregoing demonstrates that the inventions of the Asserted '705 Claims focus on specific tamperproof hardware that must interact with unique software to improve network access control technology and protect a secure computer network and the data stored thereon from infected devices. Thus, the Asserted '705 Claims are eligible as a matter of law for patent protection under step one of *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208, 216 (2014) ("*Alice*").

32. All actions and steps recited in the Asserted '705 Claims, including the act of quarantining endpoints or other computers, if necessary, requires the involvement of various hardware components running dedicated software both before, during, and after the selection and isolation of an object. Said another way, a claim directed to allowing a machine to automatically and dynamically select and isolate an unsafe device attempting to access a secure network is not simply adding a generic computer component to a fundamentally human process. Rather, it is removing the once-necessary human intervention from a fundamentally mechanical process, an "improvement in the functioning of a" networked system that simply cannot be considered directed to an abstract concept. *Enfish LLC v. Microsoft Corp.*, 822 F.3d 1327, 1339 (Fed. Cir. 2016).

33. As the specification confirms, the improvement captured by the Asserted '705 Claims is not simply quarantining an infected device, but it is instead a multi-faceted network system involving multiple interrelated software and hardware components to protect a network from known and unknown threats. Specifically, the specification of the '705 Patent discloses that to reduce the burdens of having to manually identify, connect to, isolate, and remove malicious software from an infected device, the networked system can direct an unclean computer attempting to connect to the secure network, known as the host computer, to a form of remediation, such as downloading a software patch or a software update, removing material from the host computer and/or enabling certain settings, etc. present on the host computer. (*See* Ex. 1 at 1:14–41.) Indeed, the inventions of the Asserted '705 Claims are each tethered to these advances over the art in the 2005 time frame, reciting methods and systems that automatically and dynamically detect an insecure condition by contacting a trusted computing base, receiving a response therefrom, determining whether that response contains a valid identification of cleanliness, and configuring and implementing a remediation action based on what is discovered about the state of an endpoint or "host" computer. (*See, e.g.*, *id.*, Claims 12 and 19.) More specifically, the Asserted '705 Claims require a system configured to communicate with a "trusted computing base" to determine when a response includes a valid digitally signed attestation of cleanliness, and to control access to the network accordingly. These Asserted '705 Claims are thus directed to a machine-implemented solution resolving a machine-specific problem: a machine's difficulty in detecting, isolating, and remediating infected endpoint devices (*e.g.*, host computers) to prevent contagion of and damage to the larger computer network.

34. The Asserted '705 Claims are thus directed to a machine-implemented process for (1) determining whether the host computer is required to be quarantined, (2) isolating and

inoculating the contagions (including directing the host to software programs and/or code designed to identify undesirable and/or unauthorized states) by quarantining the host, (3) limiting access to the network by the host computer so that the unsafe condition thereof can be remedied, and (4) allowing for remediation of an unsafe or infected host computer. As such, the Asserted '705 Claims recite inventions with specific applications or improvements to technologies in the marketplace and cannot be considered abstract or patent ineligible under relevant law.

35. In a recently resolved action, Citrix, a competitor of Twingate, moved a Federal Court in Florida to dismiss the claims brought against it for infringement of the '705 Patent on the grounds of patent invalidity under 35 U.S.C. 101. That motion was denied, with the court issuing a substantial opinion finding that neither *Alice* Step 1 nor *Alice* Step 2 rendered the there involved '705 Patent claims ineligible.

d. Failed IPRs of the '705 Patent

36. Fortune 100 companies accused of infringing the '705 Patent have previously filed petitions for IPRs, alleging that the claims of the '705 Patent should be held invalid as either anticipated or obvious considering art not previously considered. Ultimately, the PTAB instituted an IPR against the '705 Patent, with similar third party IPRs that were subsequently filed being joined to the first-filed and instituted IPR.

37. The PTAB eventually issued its decision holding that no claims of the '705 Patent were unpatentable, finding that no asserted prior art reference alone or in combination satisfied the limitation of "providing . . . an IP address of a quarantine server configured to serve the quarantine notification page" that was present in all claims of the '705 Patent.

38. The IPR Decision was then appealed to the CAFC, which reversed the PTAB's Decision on a few narrow procedural issues involving proof that the asserted prior art references would be combined by a person having ordinary skill in the art, as alleged by the petitioners.

39. A motion to dismiss the IPR was subsequently filed and the IPR has been dismissed with prejudice by the PTAB.

40. More recently, Netskope, Citrix and Google filed petitions seeking to have new IPR proceedings instituted against the '705 Patent. The Citrix petition was voluntarily withdrawn, while the Netskope and Google petitions were discretionarily denied.

2. The TNO Patents

41. The '717 Patent, '120 Patent, '259 Patent and '649 Patent (collectively, "the TNO Patents") were all developed at TNO, which is a prominent statutory research organization in the Netherlands focusing on applied science. Founded in 1932, TNO develops solutions to technological and social challenges facing Europe and the world, partnering with thousands of multinational corporations, universities, public sector organizations, and other subject matter experts to develop and market critical innovations. The named inventors of the TNO Patents include highly respected luminaries in their respective technical fields, including several holders of Ph.Ds.

a. The '717 Patent claims are presumed valid and patent-eligible

42. K.Mizra asserts that at least, and without limitation, Claim 1 of the '717 Patent has been directly infringed, either literally or under the doctrine of equivalents by Amazon. K.Mizra reserves the right to assert additional claims of the '717 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '717 Claims."

43. None of the Asserted '717 Claims are directed to abstract ideas, and each employs inventive concepts and are directed to patent-eligible subject matter. All claims of the '717 Patent are also presumed to be valid and enforceable against Amazon and others.

44. Indeed, the '717 Patent's specification and claims demonstrate that the need satisfied by the inventions of the Asserted '717 Claims was long-felt in the industry and thus unconventional. As one example, the '717 Patent explains that stations within wireless networks were required to transmit information before routing topology had been defined, *i.e.* when plural routes were still possible. The '717 Patent explains that this created complications because wireless stations would have to support different modes of communication while operating in low-power states, which was difficult to achieve. The '717 Patent further explains that wireless links between nodes were hindered by barriers and that using the same frequency for upward and downward link routes required synchronization, which led to further complications by requiring wireless stations to support synchronization and powerful radio frequency signals, again at low-power states. The '717 Patent thus provides "a wireless network wherein routes for communication can be determined without the need to forward information via stations for which the routes have not yet been defined," a method of operating such a network, and "stations for flexible use in wireless networks." (Ex. 2 at 3:6-15.)

45. The specification (including the provisions quoted above), the figures (including those included below), and the text related to the figures further illustrate the innovative network architecture captured by the Asserted '705 Claims. These figures include the following:

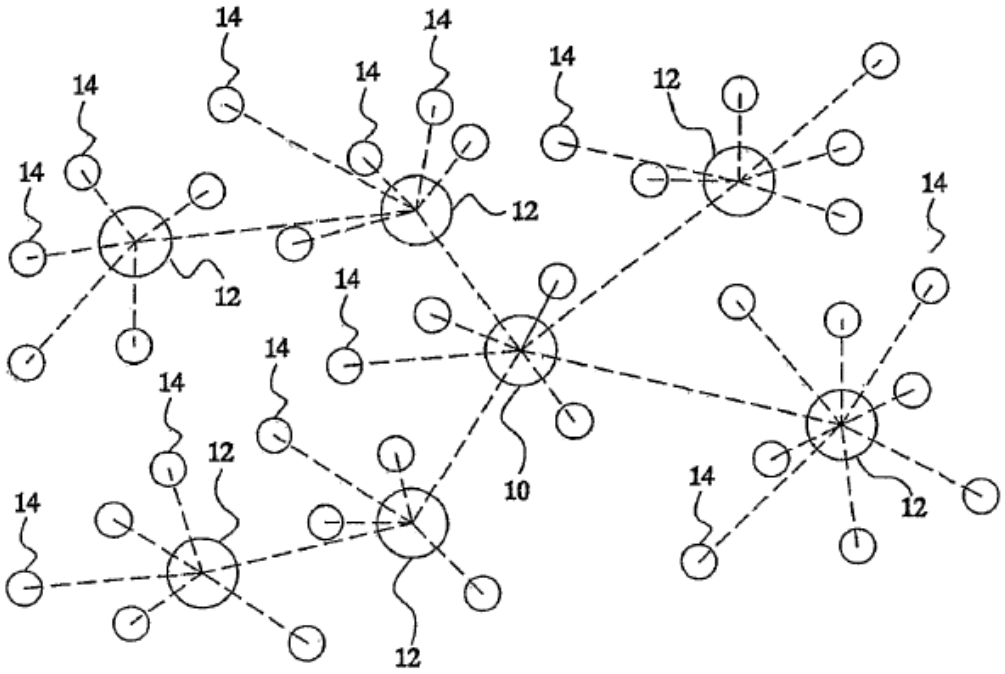

Fig.1

(*See id.* at Fig. 1.)

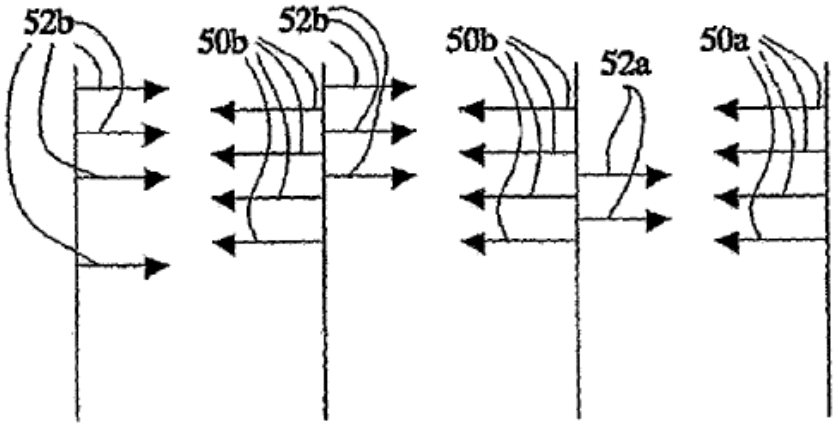

Fig.5a

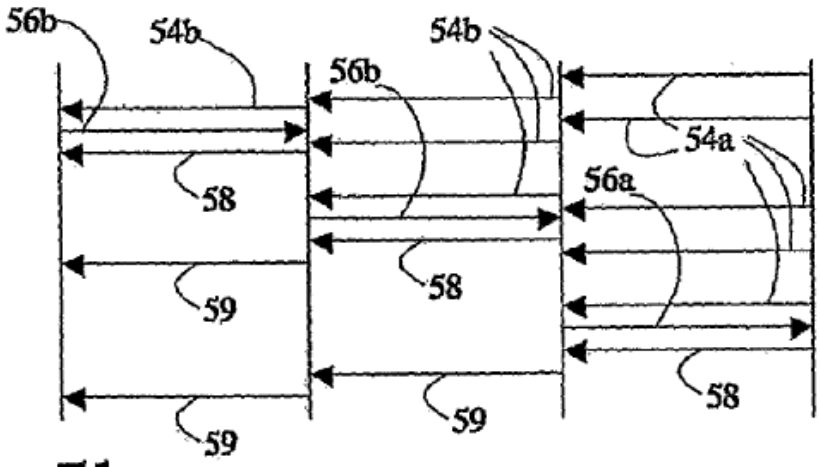

Fig.5b

(*See id.* at Fig. 5.)

46. The foregoing figures, which illustrate a wireless network and signaling protocol, demonstrate that the inventions of the Asserted '717 Claims focus on specific, highly efficient ways for a station to join a wireless network, thereby improving the operation and performance of the wireless network, as a whole. In particular, an unassociated station transmits an association request message and switches to an associated state upon reception of an association grant from

an association unit in response to one of its request messages. The association grant establishes an operating route associated with the newly associated station for exchanging operating messages with the association unit. Thereafter, the station can transmit association grants in response to received association request messages from additional stations seeking to join the network. This improves the efficiency and operability of the wireless network by facilitating the addition of new stations to the network. Thus, the Asserted '717 Claims are eligible as a matter of law for patent protection under step one of *Alice*.

47. All actions and steps recited in the Asserted '717 Claims, including the acts of transmitting association request messages, transmitting association grant messages, establishing operating routes and exchanging operating messages, and enabling stations to transmit association grant messages, require the involvement of various hardware components running dedicated software both before, during, and after the association of a station. Said another way, a claim directed to allowing a wireless network to automatically and dynamically associate a new station with the network and facilitate the addition of other stations is not simply adding a generic computer component to a fundamentally human process. Rather, it is an "improvement in the functioning of a" networked system that simply cannot be considered directed to an abstract concept. *Enfish*, 822 F.3d at 1339.

48. As the specification confirms, the improvement captured by the Asserted '717 Claims is not simply associating various devices into a wireless network, but is instead a specific network architecture whereby a particular unit operates to receive access requests and provide access grants, as well as providing operation information so that newly added wireless stations can transmit access grants to additional stations. Specifically, the specification of the '717 Patent discloses that to improve the efficiency of adding stations to a wireless network, a station seeking

to join the network operates initially in a non-associated state transmitting access requests. (*E.g.*, Ex. 2 at 8:25-32.) Subsequently, the station receives an access grant message and switches to an associated state. (*Id.* at 9:3-13.) In this state, the station establishes a dedicated communication route with the association unit or with a second station acting as an intermediary between the newly added station and the association unit, thereby ensuring efficient communication within the wireless network. (*Id.*) The newly added station can thereafter relay association grant messages to additional new stations, further enhancing the scope and efficiency of the wireless network. (*Id.* at 9:41-53.)

49. The '717 Patent specification further emphasizes that, once stations are added to the network via association grant messages, they have "well defined routes" to communicate within the network. (*Id.* at 9:63-10:6.) Thus, rather than requiring specific access routes in order to join the network initially, which creates needless complexity and power usage, the '717 Patent teaches that well defined routes are established after the station has been added to the network via an efficient method.

50. The Asserted '717 Claims are thus directed to a machine-implemented process for efficiently adding a new station to a wireless network and configuring that station to facilitate the addition of further stations to the network, improving the overall operation of the network itself. As such, the Asserted '717 Claims recite inventions with specific applications or improvements to technologies in the marketplace and cannot be considered abstract or patent ineligible under relevant law.

b. The '120 Patent claims are presumed valid and patent-eligible

51. The '120 Patent is directed to a critical advancement in machine learning, namely, an effective method of estimating hyperparameters given large potential datasets. The sole named

inventor of the '120 Patent, Stephan Raaijmakers, is a pioneer in artificial intelligence, holding a Ph.D. in Machine Learning from Tilburg University and a chair on Communicative AI at Leiden University. Dr. Raaijmakers is also the author of the book *Deep Learning for Natural Language Processing*, distributed by Simon & Schuster.

52. K.Mizra asserts that at least, and without limitation, Claim 1 of the '120 Patent has been directly infringed, either literally or under the doctrine of equivalents. K.Mizra reserves the right to assert additional claims of the '120 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '120 Claims."

53. None of the Asserted '120 Claims are directed to abstract ideas, and each employs inventive concepts and is directed to patent-eligible subject matter. All claims of the '120 Patent are also presumed to be valid and enforceable against Amazon and others.

54. Indeed, the '120 Patent's specification and claims demonstrate that the need satisfied by the inventions of the Asserted '120 Claims was long-felt in the industry and thus unconventional. The '120 Patent is directed to the use of "hyperparameters" to train machine learning systems. "Hyperparameters" are parameters that can be used to determine, for example, decision criteria to be used in classifying objects. (Ex. 3 at 1:20-23.) Hyperparameters in turn must be determined using particular methods. The '120 Patent specification references two papers discussing a method of estimating parameters called "cross-entropy" and explains why this approach is not suitable for determining hyperparameters in the machine learning context. (*Id.* at 3:44-59.) Rather, the '120 Patent describes an iterative method of determining hyperparameters by drawing a random sample of hyperparameter vectors from a set of possible hyperparameter vectors, selecting from the sample a hyperparameter vector producing the best result, updating the

estimate using the selected hyperparameter vector, and repeating the process to determine the best hyperparameters. (*Id.* at 2:1-19.)

55. The '120 Patent thus identifies the conventional method of determining parameters and explains a specific technical improvement, namely:

> The present invention is based upon the insight that methods similar to the cross-entropy method can be adapted for hyperparameter estimation by adding a best performance preservation feature. The present invention benefits from the additional insight that the best performance in a particular iteration may be preserved by restricting the sample using the best performance of the previous iterations.

(*Id.* at 6:61-67.) The '120 Patent thus specifically explains a technological solution to improving known methods of machine learning parameter estimation, and the Asserted '120 Claims directed to those improvements are patent eligible. *See, e.g.*, *Aon Re, Inc. v. Zesty.AI, Inc.*, 791 F. Supp. 3d 531, 539 (D. Del. 2025) ("Although the patent does not purport to invent machine learning algorithms, computer hardware, or the general concept of aerial imagery analysis or risk assessment, it does specifically define how these known components must be structured to achieve a supposed technological improvement in accuracy and efficiency to the combination of those components."); *cf. Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205, 1213 (Fed. Cir. 2025) (holding claims ineligible where "the claims do not delineate steps through which the machine learning technology achieves an improvement").

56. Because the Asserted '120 Claims are directed to a specific improvement in machine learning technology, the claims recite inventions with specific applications or improvements to technologies in the marketplace and cannot be considered abstract or patent ineligible under relevant law.

c. The '259 Patent claims are presumed valid and patent-eligible

57. K.Mizra asserts that at least, and without limitation, Claim 1 of the '259 Patent has been directly infringed, either literally or under the doctrine of equivalents. K.Mizra reserves the right to assert additional claims of the '259 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '259 Claims."

58. None of the Asserted '259 Claims are directed to abstract ideas, and each employs inventive concepts and is directed to patent-eligible subject matter. All claims of the '259 Patent are also presumed to be valid and enforceable against Amazon and others.

59. Indeed, the '259 Patent's specification and claims demonstrate that the need satisfied by the inventions of the Asserted '259 Claims was long-felt in the industry and thus unconventional. As one example, the '259 Patent explains that computing devices configured to detect certain sounds ("ticks") encountered problems in noisy environments where random sounds may be difficult to distinguish from sounds the devices are configured to detect. (Ex. 4 at 1:26-35.) The '259 Patent specification identifies prior art approaches involving "musical beat detection" and explains why these conventional techniques were inadequate to solve the identified problem. (*Id.* at 1:36-2:3.) The '259 Patent thus describes and claims a specific audio detection and processing system to solve this identified problem.

60. The specification (including the provisions quoted above), the figures (including those included below), and the text related to the figures further illustrate the complex audio detection and processing of the inventions captured by the Asserted '259 Claims. These figures include the following:

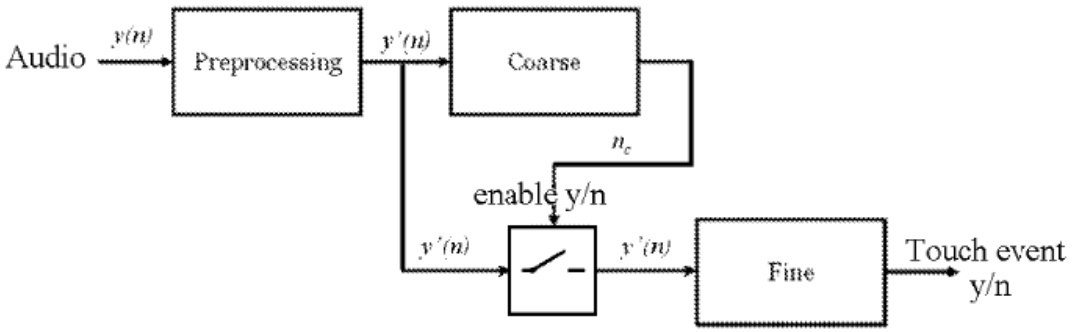

FIG. 1

(*See id.* at Fig. 1.)

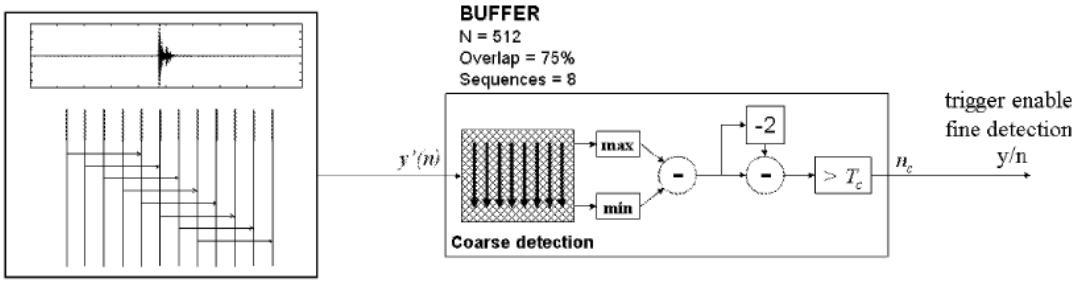

FIG. 3

(*See id.* at Fig. 3.)

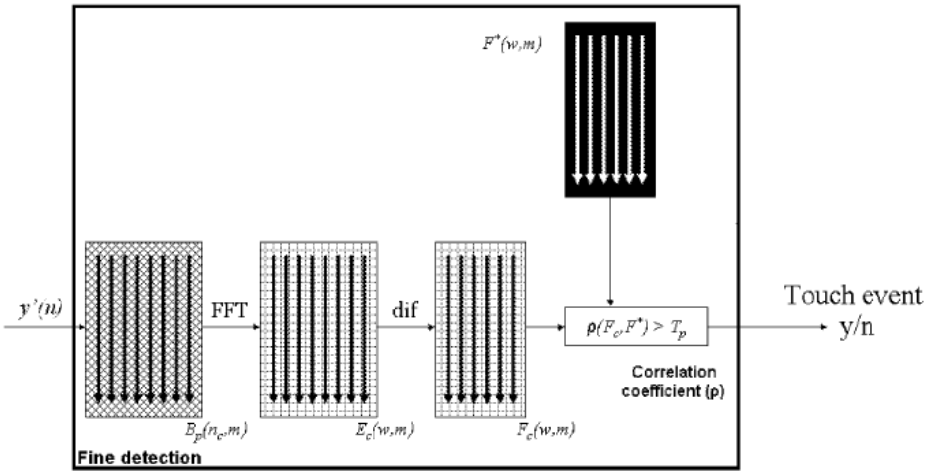

FIG. 4

(*See id.* at Fig. 4.)

61. The foregoing demonstrates that the inventions of the Asserted '259 Claims focus on specific multi-tiered hardware that must interact with unique software to conduct a specific, unconventional audio processing operation in order to properly detect an audio signal. Thus, the Asserted '259 Claims are eligible as a matter of law for patent protection under step one of *Alice*.

62. All actions and steps recited in the Asserted '259 Claims, including the acts of applying an audio signal to a course tick detection processor and enabling further processing by a fine tick detection processor, if necessary, require the involvement of various hardware components running dedicated software both before, during, and after the detection of an audio signal. Said another way, a claim directed to allowing a machine to automatically and dynamically detect a specific audio input in a noisy environment is not simply adding a generic computer component to a fundamentally human process. Rather, the claims "assert a specific and tangible improvement to computer functionality, namely the processing of audio data, that is not abstract." *Hybrid Audio, LLC v. Asus Computer Int'l*, No. 3:17-CV-05947-JD, 2019 WL 3037540, at *4 (N.D. Cal. July 11, 2019).

63. As the specification confirms, the improvement captured by the Asserted '259 Claims is not simply receiving an audio signal, but it is instead a multi-tiered processing system that first subjects a detected signal to coarse processing to determine whether further processing is necessary, and only conducts fine processing if the coarse processor determines that the detected signal is likely a desired "tick." Figure 1 above illustrates the overall system, Figure 3 illustrates the coarse processor, and Figure 4 illustrates the fine processor. (Ex. 4 at 4:13-18.) The specification provides explanations of the coarse detection operation (*id.* at 5:33-6:14) and the fine detection operation (*id.* at 6:15-7:50). The specification explains that the claimed system enables precise detection of desired audio signals while minimizing processor and power use by limiting

fine processing to signals that have first been identified as potential useful signals by a lower-power coarse processing unit. (*Id.* at 7:51-57.) The Asserted '259 Claims are thus directed to a machine-implemented solution resolving a machine-specific problem: a machine's difficulty in detecting audio signals in a noisy environment.

64. The Asserted '259 Claims are thus directed to a machine-implemented process for (1) filtering audio signals in a noisy environment to identify potential "ticks" that require further processing and input and (2) performing fine processing to definitively identify such "ticks" once the initial processing has determined that "ticks" may be present in the audio inputs. As such, the Asserted '259 Claims recite inventions with specific applications or improvements to technologies in the marketplace and cannot be considered abstract or patent ineligible under relevant law.

d. The '649 Patent claims are presumed valid and patent-eligible

65. K.Mizra asserts that at least, and without limitation, Claim 17 of the '649 Patent has been directly infringed, either literally or under the doctrine of equivalents. K.Mizra reserves the right to assert additional claims of the '649 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '649 Claims."

66. None of the Asserted '649 Claims are directed to abstract ideas, and each employs inventive concepts and is directed to patent-eligible subject matter. All claims of the '649 Patent are also presumed to be valid and enforceable against Amazon and others.

67. Indeed, the '649 Patent's specification and claims demonstrate that the need satisfied by the inventions of the Asserted '649 Claims was long-felt in the industry and thus unconventional. As one example, the '649 Patent explains that known methods of pairing electronic devices using sensory identifiers (such as audio or visual information) were inadequate because

these methods lacked a way to ensure that the devices were in fact proximate to each other. (Ex. 5 at 1:9-25.) In other words, pairing mechanisms that relied on sensory identifiers were prone to errors because sensory information could erroneously be obtained from non-proximate devices. (*Id.*) The '649 Patent thus describes and claims a specific technology to enable devices to be paired using sensory identifiers while ensuring proximity.

68. The specification (including the provisions quoted above), the figures (including those included below), and the text related to the figures further illustrate the complex audio detection and processing of the inventions captured by the Asserted '649 Claims. These figures include the following:

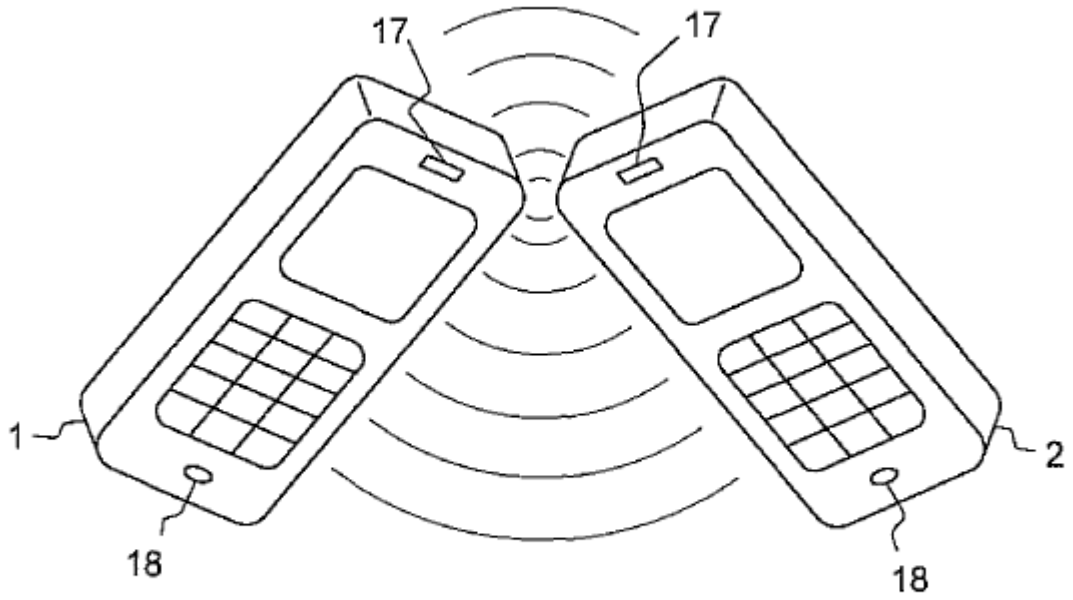

FIG. 1

(*See id.* at Fig. 1.)

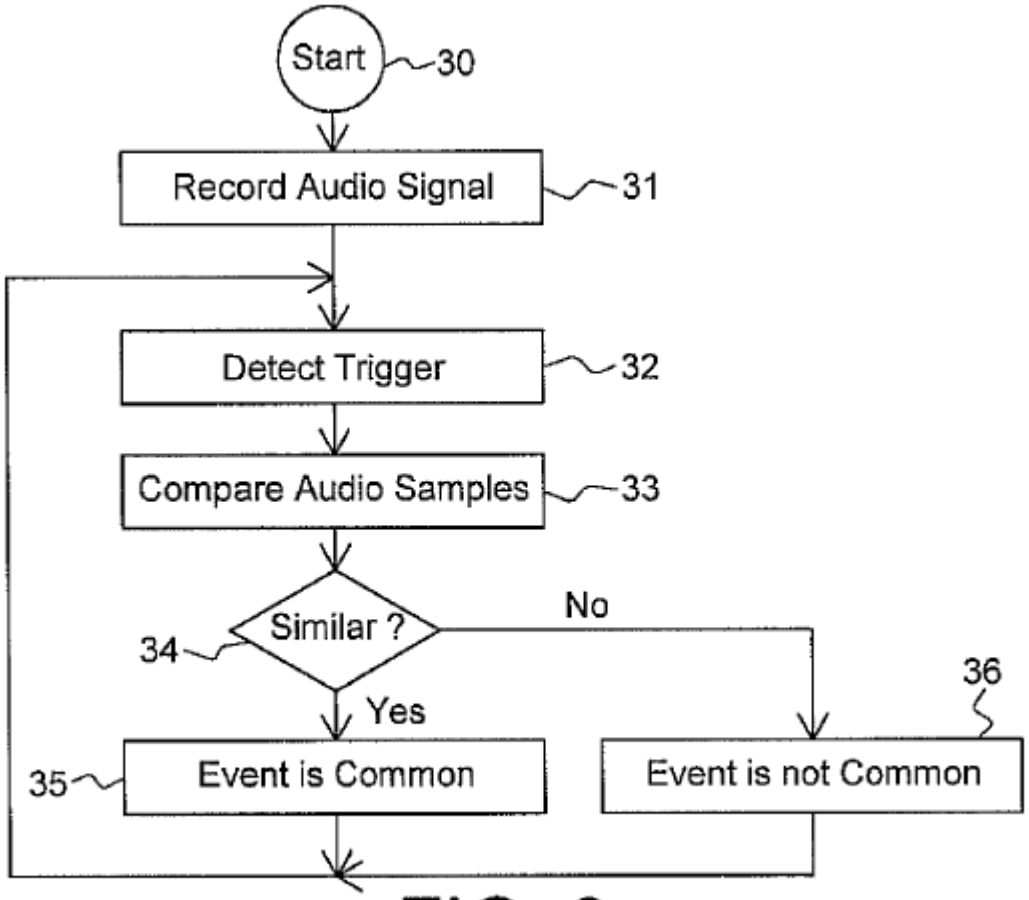

FIG. 3

(*See id.* at Fig. 3.)

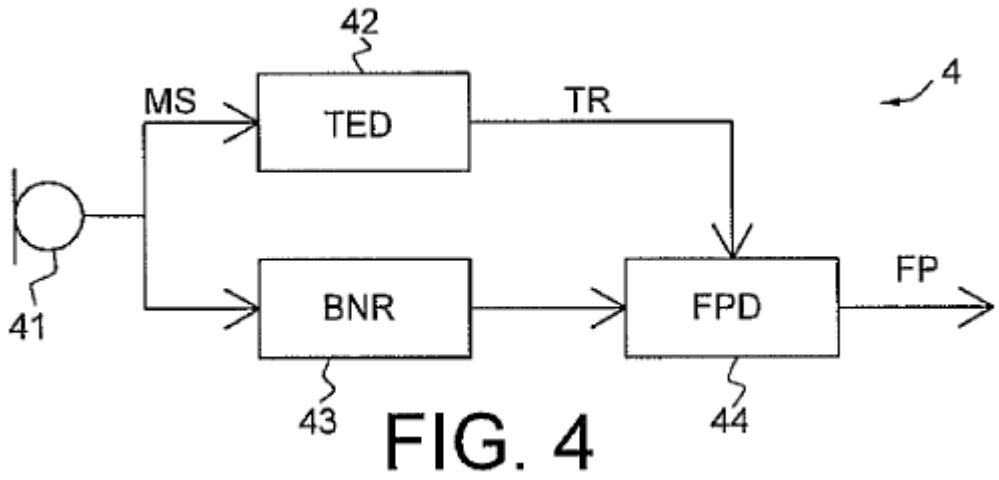

FIG. 4

(*See id.* at Fig. 4.)

69. The foregoing demonstrates that the inventions of the Asserted '649 Claims focus on a specific processing technique to ensure proximity between two devices by comparing audio samples and pairing the devices only if the comparison indicates sufficient similarity to confirm proximity. Thus, the Asserted '649 Claims are eligible as a matter of law for patent protection under step one of *Alice*.

70. All actions and steps recited in the Asserted '649 Claims, including the acts of recording audio signals in each device, detecting a sensory identifier, using the detected sensory identifier for matching the devices, selecting a common time interval of audio samples, comparing the audio samples, and deciding that the devices are proximate only if the respective audio samples are sufficiently similar, require the involvement of various hardware components running dedicated software both before, during, and after the pairing process. Said another way, a claim directed to allowing devices to automatically and dynamically detect and pair to one another only after determining that pairing conditions are satisfied is not simply adding a generic computer component to a fundamentally human process. Rather, like the '259 Patent, the '649 Patent Claims "assert a specific and tangible improvement to computer functionality, namely the processing of audio data, that is not abstract." *Hybrid Audio*, 2019 WL 3037540 at *4.

71. As the specification confirms, the improvement captured by the Asserted '649 Claims is not simply pairing devices based on a sensory identifier, but an enhanced proximity-based pairing process that avoids erroneous pairing attempts (false matches) by ensuring that devices are truly in proximity and meant to be paired. The '649 Patent specification explains that a sensory identifier that acts as a pairing trigger may be sampled by both devices, and the samples compared to ensure that the two devices detected the same sensory identifier. (Ex. 5 at 6:3-18.) The specification further explains that the pairing process may involve specific hardware

components, including a microphone, a "touch event detector," a "background noise recorder," and a "fingerprint determination unit," as shown in Figure 4 above. (*Id.* at 7:31-37.) The Asserted '649 Claims are thus directed to a machine-implemented solution resolving a machine-specific problem: reliably pairing electronic devices while avoiding false matches.

72. The Asserted '649 Claims are thus directed to a machine-implemented process for (1) recording audio signals in electronic devices, (2) detecting sensory identifiers, and (3) using the detected sensory identifiers to match the devices by comparing audio samples to ensure that the devices are in proximity. As such, the Asserted '649 Claims recite inventions with specific applications or improvements to technologies in the marketplace and cannot be considered abstract or patent ineligible under relevant law.

3. The '282 Patent

73. The '282 Patent was initially assigned to Hammerhead Systems, Inc., who thereafter assigned them to Brixham Solutions Ltd. Brixham Solutions Ltd. assigned the '282 Patent to Global Innovation Aggregators LLC, who in turn assigned it to K.Mizra. Hammerhead Systems was a networking company that dissolved in 2009.

74. K.Mizra asserts that at least, and without limitation, Claim 1 of the '282 Patent has been directly infringed, either literally or under the doctrine of equivalents. K.Mizra reserves the right to assert additional claims of the '282 Patent, including both independent and dependent claims, pursuant to the Court's (and other applicable) rules and procedures and as discovery progresses. These claims are referred to herein as the "Asserted '282 Claims."

75. The '282 Patent teaches novel systems and methods for improving scalability and efficiency in handling large volumes of events occurring in network and network management tasks. The '282 Patent's claimed systems and methods combine specific hardware and software

components in unconventional ways. Conventional network management systems (NMS) systems faced several challenges as the size of data communications networks grew in size. As networks expanded, NMS's required increased processing power to manage network performance, identify and fix problems or malfunctions in the network, and plan for further network growth without becoming overwhelmed by the number of events occurring in the network that are generated by network nodes and elements.

76. Through its novel technological innovation, the '282 Patent provides systems and methods for a scalable and robust NMS architecture that can efficiently manage increasing network demands by implementing server clusters, load balancing, and distributed gateway and adapter systems. This allows for incremental addition of network management resources without interrupting network operations, thereby optimizing performance and scalability.

77. The '282 Patent is directed to improvements to computer network technology, more particularly, "a network management system that is robust, scalable, and capable of efficiently handling a growing network." (Ex. 6 at 1:23-25.) A network management system "refers to a system responsible for managing a network, and facilitates the communication between a carrier's OSS and NEs in the network." (*Id.* at 2:37-40.) Prior network management systems ("NMSs") "face[d] many challenges as networks grow in size[,]" including the "thousands of events [that] can be generated when a fault is discovered[,]" resulting in it "easily be[ing] overwhelmed." (*Id.* at 1:17-22.) The '282 Patent specification notes that "[i]t would be desirable to develop a network management system that is robust, scalable, and capable of efficiently handling a growing network." (*Id.* at 1:22-25.) The problem the inventors sought to address was thus rooted in computer networking technology, *i.e.*, the technical limitations of prior art NMSs in their scalability, efficiency, and robustness. (*See id.*) This problem—NMSs with limited scalability,

efficiency, and robustness—is not a problem that arises in longstanding human activity, but is rather rooted in machine-based networking.

78. Recognizing this problem, the inventors provided a "scalable network management system[,]" (*id.* at 2:12-13), with, *e.g.*, a distributed architecture. Specifically, the inventors achieved, *e.g.*, "[h]igh scalability . . . by using server clustering and distributed adapter architecture, with the load balancing scheme to connect the server cluster and NB and SB adapters in northbound and southbound directions respectively." (*Id.* at 7:65-8:2.) In this distributed architecture, "multiple instances of such application servers can run simultaneously to distribute a workload and protect each other [where] [e]ach server instance by itself is a complete functioning NMS server that runs independently[.]" (*Id.* at 8:12-15.) An example of this is shown below:

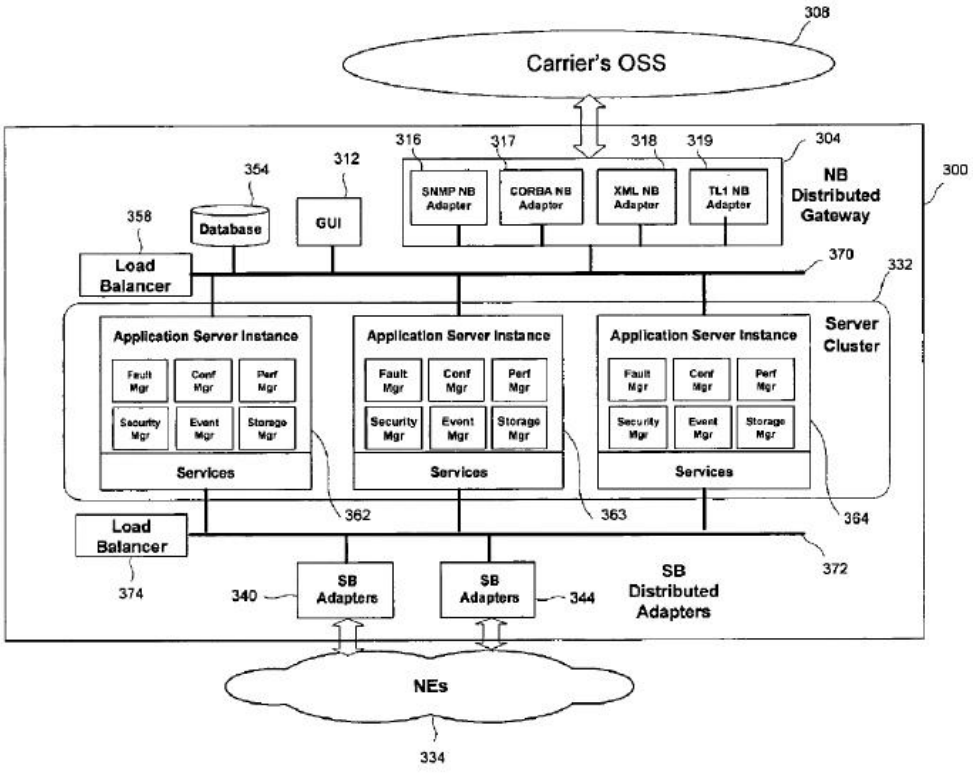

FIG. 3A: NB distributed gateway, SB distributed adapters, server cluster, & load balancers

(*Id.* at Fig. 3A.) The disclosed distributed architecture provides a specific, technical solution to a technical problem and tangible technical benefits rather than abstract ideas, including the ability

to increase NMS processing power, efficiency, and flexibility. (*Id.* at 9:47-60.) The '282 Patent notes that other technical benefits also flow from embodiments of the improved distributed architecture, including:

- "With this architecture of an NB gateway decoupled from NMS application servers, multiple NB gateways, each of which includes a certain number of NB adapters, can be designed and distributed on different workstations for scalability and performance optimization." (*Id.* at 9:39-44.)
- "An instance of the server can also be shut down for any reason without interrupting the server functionality as a whole. When a server is shutdown, its associated software modules (NB gate way 304, GUI312, or SB adapters 340 or 344) automatically re-establish the association with another server instance (one of existing servers 362-364 or a standby server) based on certain criterion, such as selecting the lightest loaded server." (*Id.* at 9:58-65.)
- "[T]he distributed SB adapter-based NMS architecture can Support a flexible network partitioning to manage different management domains based on the amount of management traffic in each domain in a large scale network." (*Id.* at 9:67-10:2.)

The improvement over and difference from prior conventional NMSs is a technical one: an NMS with the improved distributed architecture can be enabled to re-establish an association with another server to ensure functionality during an interruption. Thus, the Asserted '282 Claims are eligible as a matter of law for patent protection under step one of *Alice*.

79. All actions and steps recited in the Asserted '282 Claims, including the acts of receiving, at a first application server instance selected from a plurality of application server

instances based on a load balancing process, first adapter processed information from a first adapter; processing, by the first application server instance, the first adapter processed information based on an event management service to produce application processed information; sending, by the first application server instance, the application processed information to a gateway device, wherein the gateway device is one of a plurality of gateway devices respectively associated with the plurality of application server instances and is configured to transfer the application processed information to a second adapter of a plurality of second adapters configured to process the application processed information and transfer the second adapter processed information to an operation support system device; and in response to determining that the first application server instance has become disabled, facilitating establishing an association between the first adapter and a second application server instance of the plurality of application server instances and between the gateway device and the second application server instance require the involvement of various hardware components running dedicated software both before, during, and after the establishment of an association with a disabled server instance. Said another way, a claim directed to allowing a machine to automatically and dynamically manage associations within a network to ensure that the network operates properly and efficiently with minimal interruptions is not simply adding a generic computer component to a fundamentally human process. Rather, it is a fundamentally mechanical process, an "improvement in the functioning of a" networked system that simply cannot be considered directed to an abstract concept. *Enfish LLC v. Microsoft Corp.*, 822 F.3d 1327, 1339 (Fed. Cir. 2016).

80. During patent examination, the examiner allowed claim 1 because the closest prior art references "do not sufficiently teach or disclose all of the recited limitations of the amended independent claims, and claim 1 in particular, including the recited feature of 'in response to

determining that the first application server instance has become disabled, facilitating establishing an association between the first adapter and a second application server instance of the plurality of application server instances and between the gateway device and the second application server instance …." (Ex. 7 (3/17/2014 Notice of Allowance, '282 Patent File History) at 2.) That the Patent Office found the closest prior art did not disclose these additional inventive elements supports their being non-routine and unconventional. Accordingly, claim 1 of the '282 Patent recites additional elements that were not "well-understood, routine, or conventional," which provide "something more" sufficient to transform any alleged abstract idea into eligible subject matter.

B. Amazon's Accused Instrumentalities And Services

81. Amazon is not and has never been a licensee of any of the Asserted Patents nor had or has any rights to use technologies covered by the Asserted Patents. Amazon thus has no ownership or other rights (and is entitled to no rights) relating to the Asserted Patents.

82. Amazon has been making, selling, using, and offering for sale various products and services that infringe the Asserted Patents in violation of 35 U.S.C. § 271. These include, but are not limited to, the following:

- AWS Verified Access ('705 Patent);
- Thread-enabled Amazon devices, including without limitation Thread-enabled Amazon Echo devices ('717 Patent);[1]
- Amazon SageMaker AI ('120 Patent);
- Alexa Emergency Assist ('259 Patent);

[1] The Thread Specification is incorporated by reference.

- Amazon Echo devices incorporating Echo Spatial Perception ('649 Patent); and
- Amazon CloudWatch ('282 Patent)

(collectively, the "Accused Instrumentalities"). The sale, offer for sale, use and/or manufacture in the United States of the Accused Instrumentalities constitutes infringement of at least and without limitation, the asserted claims directly, either literally or under the doctrine of equivalents.

C. K.Mizra's Efforts To Work With Amazon

83. K.Mizra has contacted Amazon on several occasions seeking to discuss Amazon's infringement. For example, K.Mizra contacted Amazon at least as early as January 6, 2025 in a letter addressing Amazon's potential infringement of the '705 Patent, including a claim chart showing Amazon's infringement of the '705 Patent. (*See* Ex. 8.)

84. Thereafter, counsel for K.Mizra communicated with Amazon to resolve this dispute, but no agreement has been reached.

COUNT I
(Patent Infringement under 35 U.S.C. § 271 of the '705 Patent)

85. K.Mizra incorporates paragraphs 1 through 84 as though fully set forth herein.

86. The '705 Patent includes 19 claims.

87. Amazon has directly infringed one or more claims of the '705 Patent by making, importing, using, offering for sale, and/or selling AWS Verified Access, all in violation of 35 U.S.C. § 271(a).

88. Based on publicly available information, AWS Verified Access satisfies every element of at least Claim 19 of the '705 Patent.

89. Claim 19 of the '705 Patent states:

[preamble] A computer program product for protecting a network, the computer program product being embodied in a non-transitory computer readable medium and comprising computer instructions for:

[A] detecting an insecure condition on a first host that has connected or is attempting to connect to a protected network,

[B] wherein detecting the insecure condition includes:

[B1] contacting a trusted computing base associated with a trusted platform module within the first host,

[B2] receiving a response, and determining whether the response includes a valid digitally signed attestation of cleanliness,

[C] wherein the valid digitally signed attestation of cleanliness includes at least one of an attestation that the trusted computing base has ascertained that the first host is not infested, and an attestation that the trusted computing base has ascertained the presence of a patch or a patch level associated with a software component on the first host;

[D] when it is determined that the response does not include a valid digitally signed attestation of cleanliness, quarantining the first host, including by preventing the first host from sending data to one or more other hosts associated with the protected network,

[E] wherein preventing the first host from sending data to one or more other hosts associated with the protected network includes

[E1] receiving a service request sent by the first host, serving a quarantine notification page to the first host when the service request comprises a web server request,

[E2] and in the event the service request comprises a DNS query, providing in response an IP address of a quarantine server configured to serve the quarantine notification page if a host name that is the subject of the DNS query is not associated with a remediation host configured to provide data usable to remedy the insecure condition; and

[F] permitting the first host to communicate with the remediation host.

(Ex. 1 at 22:14-49.)

90. As for the preamble of Claim 19, to the extent that it is determined to be limiting, AWS Verified Access provides the features described in the preamble, which recites a "computer program product for protecting a network." Amazon's AWS Verified Access provides for secure access to protected network resources.

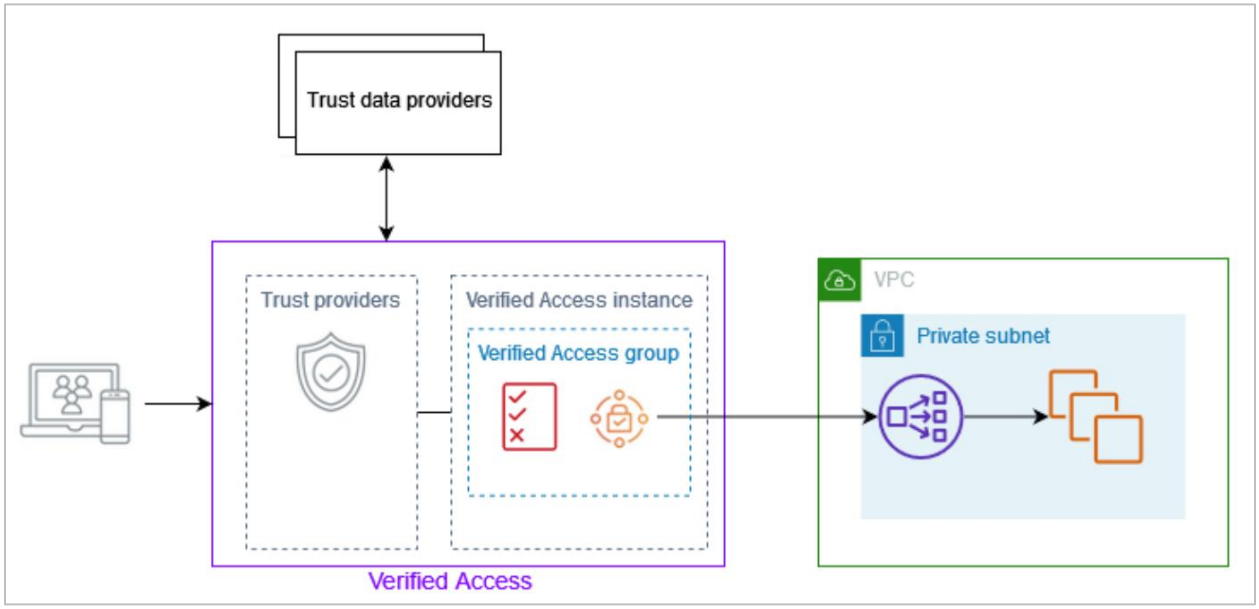

(*See* How to use AWS Verified Access logs to write and troubleshoot access policies, available at https://aws.amazon.com/blogs/security/how-to-use-aws-verified-access-logs-to-write-and-troubleshoot-access-policies/ (last accessed January 2026 and incorporated by reference).) For example, Amazon touts that its AWS Verified Access product "enables secure access to applications" and "evaluates each application request and gives users access to each application only when they meet the specified security requirements." (*Id.*) Amazon offers its AWS Verified Access service for sale as follows:

Pricing table for HTTP(S) applications

There are two pricing dimensions that determine your Verified Access bill for HTTP(S) applications 1) hourly charges for associated applications, and 2) per-GB charges for data processed.

Application hours

You are charged per hour for each application associated with an active HTTP(S) Verified Access endpoint. Each partial application hour (app-hour) consumed is billed as a full hour. For example, if you associate 10 applications with 10 Verified Access endpoints, each for an hour, you will be charged 10 app-hours.

GB of data processed

You are also charged per GB for data processed by Verified Access. When users request access or connect to your applications, Verified Access processes the data flowing between users and the application. For example, if all your applications using Verified Access require 10 GB of total data processing, you will be charged for that 10 GB of data processed.

Tiered pricing for app hours	Price
1 to 148,800 app hours (covers up to 200 apps in a month)	$0.27/hr
148,800+ app hours	$0.20/hr
Note: 200 apps associated for 31 days equals 148,800 app hours	
Per GB of data processed	$0.02

Example 2:

You created a Verified Access in US East (Ohio) and associated 300 applications. These 300 applications were associated for a month with 31 days. This will result in 223,200 app hours. Also, each application required 1 GB for data processing. The total charge will be $55,062.00, and the below table provides the details.

First 148,800 app hours	148,800 hours * $0.27	$40,176.00
Next 74,400 (223,200-148,800) app hours	74,400 hours * $0.20	$14,880.00
Data processing	1 GB * 300 applications * $0.02	$6.00
Total charge		**$55,062.00**

(*See* AWS Verified Access Pricing, available at https://aws.amazon.com/verified-access/pricing/) (last accessed January 2026 and incorporated by reference).) Accordingly, and to the extent that the preamble of Claim 19 is somehow limiting, AWS Verified Access would meet the limitation.

91. Limitation A of Claim 19 requires "detecting an insecure condition on a first host that has connected or is attempting to connect to a protected network." AWS Verified Access also meets all the requirements of limitation A of Claim 19. Amazon's AWS Verified Access "enhances [owners'] security posture by allowing [owners] to define fine-grained access policies based on a user's identity and device security state, and enforcing policies on every access request."

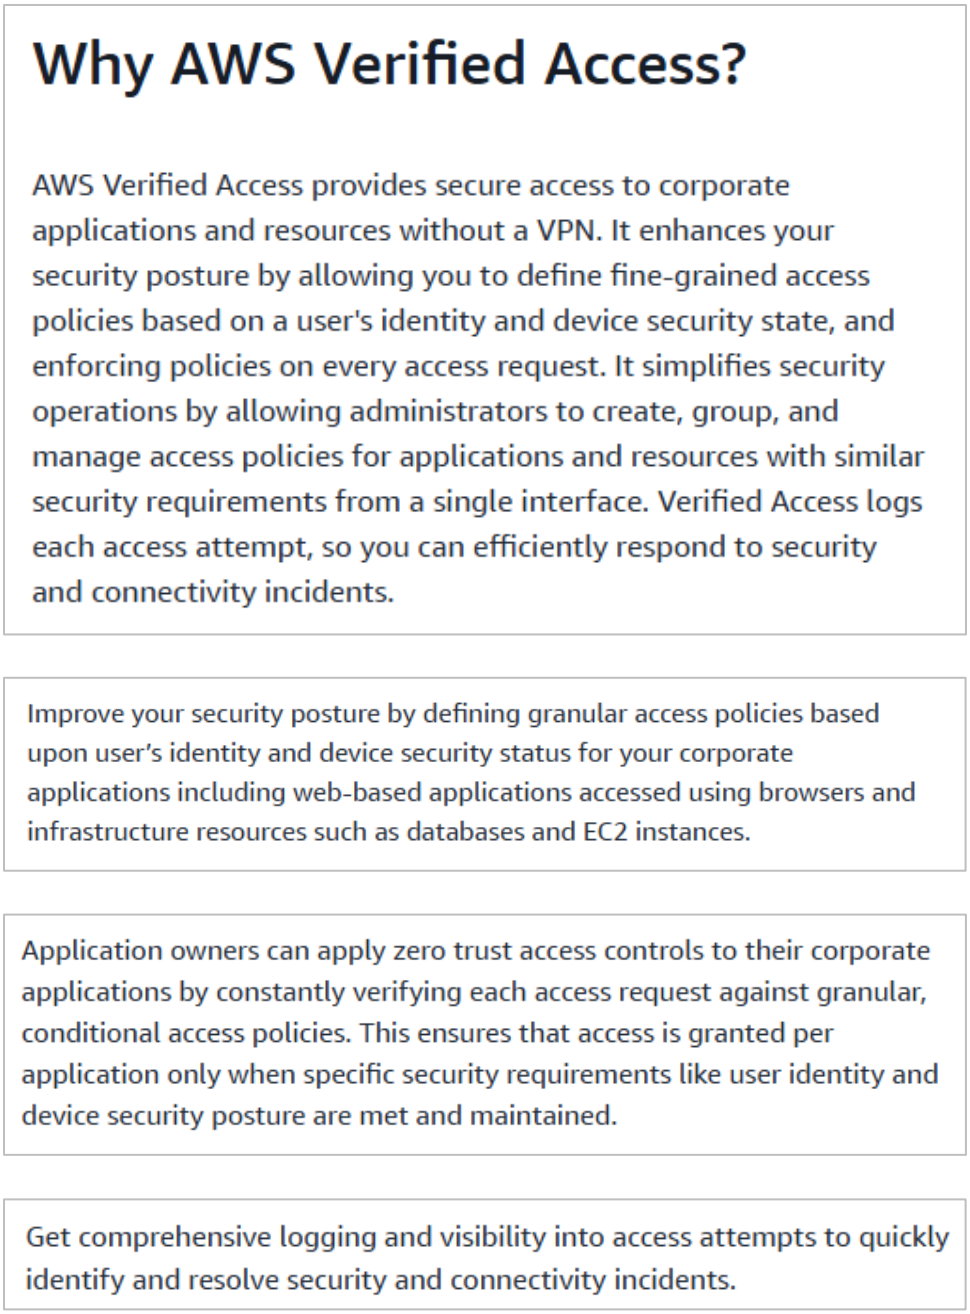

(Secure Remote Access – AWS Verified Access – AWS, available at https://aws.amazon.com/verified-access/ (last accessed January 2026 and incorporated by reference).) AWS Verified Access allows owners to "apply zero trust access controls to their corporate applications by constantly verifying each access request against granular, conditional access policies," which "ensures that access is granted per application only when specific security

requirements like user identity and device security posture are met and maintained. (*Id.*) AWS Verified Access supports device trust providers such as Jamf, CrowdStrike and JumpCloud.

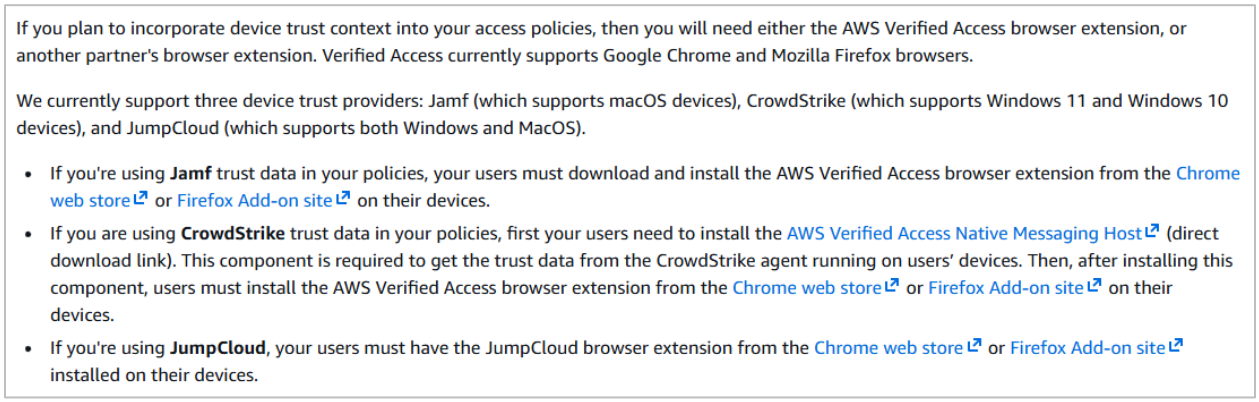

If you plan to incorporate device trust context into your access policies, then you will need either the AWS Verified Access browser extension, or another partner's browser extension. Verified Access currently supports Google Chrome and Mozilla Firefox browsers.

We currently support three device trust providers: Jamf (which supports macOS devices), CrowdStrike (which supports Windows 11 and Windows 10 devices), and JumpCloud (which supports both Windows and MacOS).

- If you're using **Jamf** trust data in your policies, your users must download and install the AWS Verified Access browser extension from the Chrome web store or Firefox Add-on site on their devices.
- If you are using **CrowdStrike** trust data in your policies, first your users need to install the AWS Verified Access Native Messaging Host (direct download link). This component is required to get the trust data from the CrowdStrike agent running on users' devices. Then, after installing this component, users must install the AWS Verified Access browser extension from the Chrome web store or Firefox Add-on site on their devices.
- If you're using **JumpCloud**, your users must have the JumpCloud browser extension from the Chrome web store or Firefox Add-on site installed on their devices.

(Third-party trust provider context for Verified Access trust data, available at https://docs.aws.amazon.com/verified-access/latest/ug/trust-data-third-party-trust.html (last accessed January 2026 and incorporated by reference).)

92. Limitation B1 of Claim 19 requires that "detecting [an] insecure condition includes . . . contacting a trusted computing base associated with a trusted platform module within the first host." AWS Verified Access meets these requirements by, for example, accessing data or security applications located on the device attempting to access a network resource. For example, using JumpCloud, AWS Verified Access accesses a "Secure Enclave" on a macOS device or a "Trusted Platform Module (TPM)" on a Windows or Linux device.

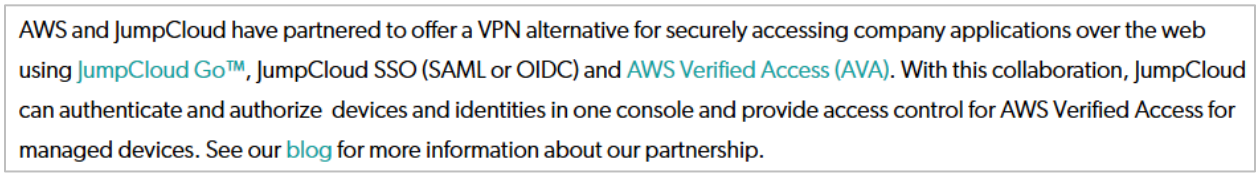

AWS and JumpCloud have partnered to offer a VPN alternative for securely accessing company applications over the web using JumpCloud Go™, JumpCloud SSO (SAML or OIDC) and AWS Verified Access (AVA). With this collaboration, JumpCloud can authenticate and authorize devices and identities in one console and provide access control for AWS Verified Access for managed devices. See our blog for more information about our partnership.

(Integrate with AWS Verified Access, available at https://jumpcloud.com/support/integrate-with-aws-verified-access (last accessed January 2026 and incorporated by reference).)

JumpCloud Go supports the following device types that meet these hardware requirements:

- macOS devices with a Secure Enclave.
- Windows devices with a Trusted Platform Module (TPM) 2.0.
- Linux devices (GNOME-based distros) with a Trusted Platform Module (TPM) 2.0.

The JumpCloud agent has to be installed and running on macOS, Windows, and Linux devices. See Install the JumpCloud Agent.

(Get Started: JumpCloud Go™, available at https://jumpcloud.com/support/get-started-jumpcloud-go (last accessed January 2026 and incorporated by reference).)

93. Limitation B2 of Claim 19 requires that "detecting the insecure condition" also includes "receiving a response and determining whether the response includes a valid digitally signed attestation of cleanliness." AWS Verified Access also meets all the requirements of limitation B2. For example, when a remote device ("first host") first initiates a request to access corporate resources (*i.e.*, a protected network), AWS Verified Access (in conjunction with the JumpCloud service), the device undergoes an enrolment process whereby it is assigned an ID Token. Thereafter, a DUK (Device User Key) is created for the device, which is added to a "JWT (JSON Web Token) that is signed by the JumpCloud Agent certificate unique to the device and used by the JumpCloud Agent to authenticate all traffic back to JumpCloud."

JumpCloud Go anchors its trust to both the device and the user. During the enrollment process the user authenticates to JumpCloud's OpenID Connect service and is issued an ID Token for the JumpCloud Go application. A new DUK (Device User Key) is created from the Secure Enclave (on Apple devices) or the Trusted Platform Module (TPM) (on Windows systems). This is added to a JWT (JSON Web Token) that is then signed by the JumpCloud Agent certificate unique to the device and used by the JumpCloud Agent to authenticate all traffic back to JumpCloud.

(JumpCloud Go™, available at https://jumpcloud.com/wp-content/uploads/2023/11/202310-WP-JumpCloudGo.pdf (last accessed January 2026 and incorporated by reference).) The JWT's signature is subject to a multi-step validation and inspection process:

> The JumpCloud Go enrollment request is validated by JumpCloud's DURT (Device User Refresh Token) API through the following process:
>
> 1. The JWT's signature is validated by first ensuring the certificate in the x5c claim in the JWT header is a valid JumpCloud Agent certificate in good standing.
> 2. The JWT signature is checked to ensure it was signed by the private key of the same agent certificate.
> 3. The JWT is inspected to ensure it hasn't expired and is otherwise valid.
> 4. The embedded ID Token is validated back to the JumpCloud OpenID Connect service via its embedded signature using the OIDC JWKS.

(*Id.*)

94. Limitation C of Claim 19 requires that "the valid digitally signed attestation of cleanliness includes at least one of an attestation that the trusted computing base has ascertained that the first host is not infested, and an attestation that the trusted computing base has ascertained the presence of a patch, or a patch level associated with a software component on the first host." AWS Verified Access meets these requirements. For example, AWS Verified Access utilizes a "trust provider," which "is a service that sends information about users and devices to AWS Verified Access, including "device information such as installed security patches or anti-virus software version."

> A trust provider is a service that sends information about users and devices to AWS Verified Access. This information is called trust context. It can include attributes based on user identity, such as an email address or membership in the "sales" organization, or device information such as installed security patches or anti-virus software version.

(Trust providers for Verified Access, available at https://docs.aws.amazon.com/verified-access/latest/ug/trust-providers.html (last accessed January 2026 and incorporated by reference).)

95. Limitation D of Claim 19 requires that "when it is determined that the response does not include a valid digitally signed attestation of cleanliness, quarantining the first host,

including by preventing the first host from sending data to one or more other hosts associated with the protected network." AWS Verified Access meets these requirements by, *e.g.*, denying all application requests that do not meet the specified security requirements.

> When a user tries to access an application, Verified Access gets their data from the trust provider and evaluates it against the policies that you set for the application. Verified Access grants access to the requested application only if the user meets your specified security requirements. All application requests are denied by default, until a policy is defined.
>
> In addition, Verified Access logs every access attempt, to help you respond quickly to security incidents and audit requests.

> **Verified Access instances** – An instance evaluates application requests and grants access only when your security requirements are met.
>
> **Verified Access endpoints** – Each endpoint represents an application. In the diagram above, the application is hosted on EC2 instances that are targets of a load balancer.
>
> **Verified Access group** – A collection of Verified Access endpoints. We recommend that you group the endpoints for applications with similar security requirements to simplify policy administration. For example, you can group the endpoints for all your sales applications together.
>
> **Access policies** – A set of user-defined rules that determine whether to allow or deny access to an application. You can specify a combination of factors, including user identity and device security state. You create a group access policy for each Verified Access group, which is inherited by all endpoints in the group. You can optionally create application-specific policies and attach them to specific endpoints.
>
> **Trust providers** – A service that manages user identities or device security state. Verified Access works with both AWS and third-party trust providers. You must attach at least one trust provider to each Verified Access instance. You can attach a single identity trust provider and multiple device trust providers to each Verified Access instance.
>
> **Trust data** – The security-related data for users or devices that your trust provider sends to Verified Access. Also referred to as *user claims* or *trust context*. For example, the email address of a user or the operating system version of a device. Verified Access evaluates this data against your access policies when it receives each request to access an application.

(How Verified Access Works, available at https://docs.aws.amazon.com/verified-access/latest/ug/how-it-works.html (last accessed January 2026 and incorporated by reference).)

96. Limitation E1 of Claim 19 requires that "preventing the first host from sending data to one or more other hosts associated with the protected network includes . . . receiving a service request sent by the first host [and] serving a quarantine notification page to the first host when the service request comprises a web server request." AWS Verified Access meets these requirements, *e.g.*, because a user attempting to access an application protected by AWS Verified Access that does not meet the security requirements will receive a 403 Unauthorized error message.

> Use a web browser with the Verified Access extension to test access to the application. Before signing in, you should be redirected to the Okta sign-in page. After successful authentication, if the device does not have a ZTA score above the threshold in the group policy, then the user receives a 403 Unauthorized error.

(Integrating AWS Verified Access with device trust providers, available at https://aws.amazon.com/blogs/networking-and-content-delivery/integrating-aws-verified-access-with-device-trust-providers/ (last accessed January 2026 and incorporated by reference).)

97. Limitation E2 of Claim 19 requires that "preventing the first host from sending data to one or more other hosts associated with the protected network includes" "in the event the service request comprises a DNS query, providing in response an IP address of a quarantine server configured to serve the quarantine notification page if a host name that is the subject of the DNS query is not associated with a remediation host configured to provide data usable to remedy the insecure condition." AWS Verified Access meets all the requirements of limitation E2 of Claim 19. For example, in connection with a 403 Unauthorized error message, AWS Verified Access provides an IP address of a quarantine server configured to serve the quarantine notification page (*e.g.*, a 403 Unauthorized error message). Accordingly, AWS Verified Access meets limitation E2 of Claim 19.

98. Limitation F of Claim 19 requires "permitting the first host to communicate with the remediation host." AWS Verified Access meets this limitation, for example, by permitting communication with "AVA Trust Providers" that act as remediation hosts.

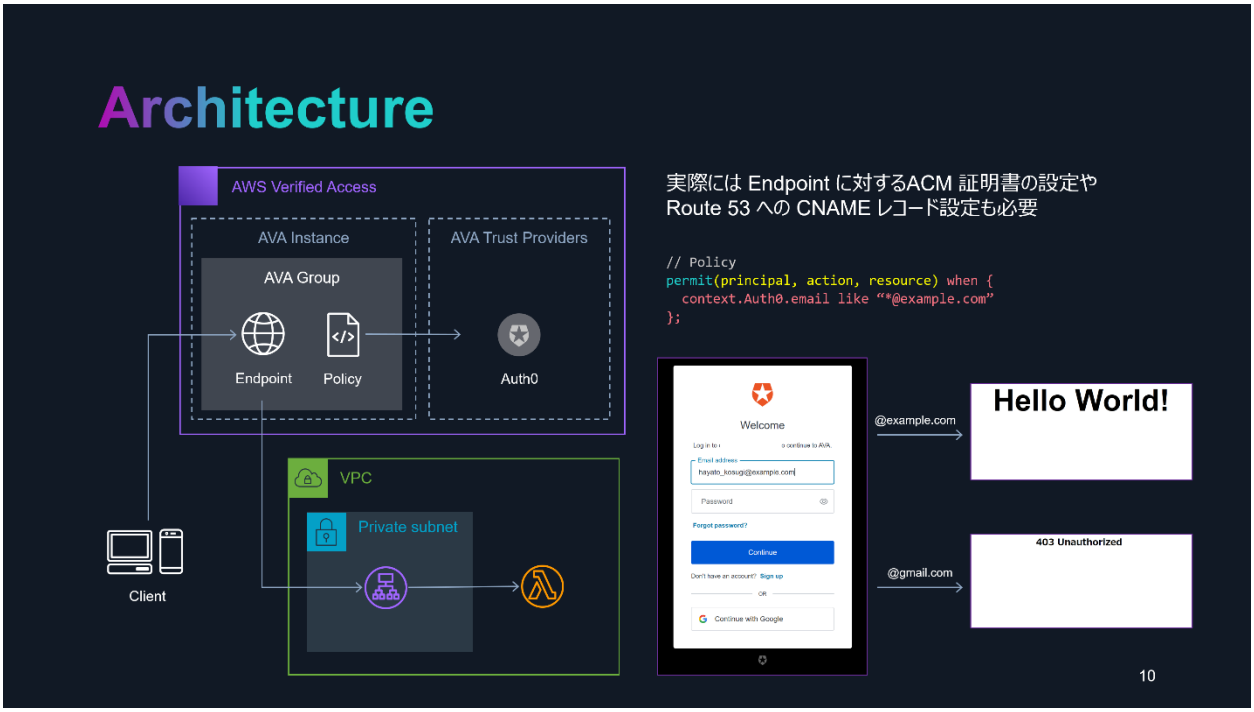

(*See* https://qiita.com/hayao_k/items/5f3b8a88a0ea75828f95 (last accessed January 2026 and incorporated by reference).) Accordingly, AWS Verified Access meets limitation F of Claim 19.

99. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '705 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of AWS Verified Access to directly infringe one or more claims of the '705 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '705 Patent and its infringement thereof no later than January 6, 2025, when it received the letter attached as Exhibit 8, and (b) Amazon intentionally causes, urges, or encourages users of AWS Verified Access to take action that, when taken, directly infringe one or more claims of the '705 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use AWS Verified Access, including infringing uses thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of AWS Verified

Access to directly infringe one or more claims of the '705 Patent, and users thereof directly infringe one or more claims of the '705 Patent. For instance, at a minimum, Amazon has supplied and continues to supply AWS Verified Access to customers while knowing that use thereof will infringe one or more claims of the '705 Patent.

100. Due to Amazon's knowledge of the '705 Patent and of its infringement thereof no later than January 6, 2025, Amazon's infringement of the '705 Patent has been and continues to be willful since no later than that date.

101. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

102. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

COUNT II
(Patent Infringement under 35 U.S.C. § 271 of the '717 Patent)

103. K.Mizra incorporates paragraphs 1 through 102 as though fully set forth herein.

104. The '717 Patent includes 20 claims.

105. Amazon has directly infringed one or more claims of the '717 Patent at least by using Thread-enabled Amazon devices in a Thread-enabled network, including without limitation Thread-enabled Amazon Echo devices ("the '717 Accused Instrumentalities") in violation of 35 U.S.C. § 271(a).

106. Based on publicly available information, the '717 Accused Instrumentalities satisfies every element of at least Claim 1 of the '717 Patent.

107. Claim 1 of the '717 Patent states:

> [preamble] A wireless communication network that comprises
>
> [A] an association unit and
>
> [B] a plurality of stations,
>
> [C] each station being configured to start up in a not-associated state and to transmit association request messages in the not-associated state,
>
> [D] each station being configured to switch to an associated state upon reception of an association grant in response to one of its association request messages,
>
> [E] the association request message and/or the association grant in response thereto establishing an operating route associated with the station for exchanging operating messages with the association unit during subsequent operation,
>
> [F] the operating route associated with the station running through a source of the association grant in response to said one of the association request messages and,
>
> [G] if the source of the association grant is not the association unit, the operating route between the source of the association grant and the association unit previously established for the source of the association grant in response to association request messages and/or the association grants for intermediate stations, in which network
>
> [H] the association unit is configured to transmit association grants as responses to received ones of the association request messages;
>
> [I] at least part of the stations is configured to become active to transmit association grants as responses to received ones of the association request messages, but only after switching to the associated state upon an association grant from the association unit or one of the at least part of the stations switched to the associated state directly or indirectly based on an association grant from the association unit.

(Ex. 2 at 14:28-52.)

108. As for the preamble of Claim 1, to the extent that it is determined to be limiting, the '717 Accused Instrumentalities provide the features described in the preamble, which recites a wireless communication network. A Thread-enabled device is configured to operate within a wireless communication network.

109. Limitation A of Claim 1 requires an association unit. A Thread-enabled wireless network will include a "Leader" responsible for managing router ID assignment and acting as a central arbiter of the network configuration state. The Leader is an association unit.

110. Limitation B of Claim 1 requires a plurality of stations. The '717 Accused Instrumentalities are "Thread Devices" (IPv6-capable devices having at least one Thread Interface and implementing the Thread Specification) and thus stations.

111. Limitation C of Claim 1 requires "each station being configured to start up in a not-associated state and to transmit association request messages in the not-associated state." The '717 Accused Instrumentalities meet this limitation by, *e.g.*, being configured to start up in a "Detached" state according to the Thread Specification and transmit "Parent Request" signals according to the Thread Specification.

112. Limitation D of Claim 1 requires "each station being configured to switch to an associated state upon reception of an association grant in response to one of its association request messages." The '717 Accused Instrumentalities meet this limitation by, *e.g.*, switching to an "Attached" state according to the Thread Specification upon receipt of a "Child ID Response" signal from the Leader according to the Thread Specification.

113. Limitation E of Claim 1 requires "the association request message and/or the association grant in response thereto establishing an operating route associated with the station for exchanging operating messages with the association unit during subsequent operation." Use of the

'717 Accused Instrumentalities in a Thread-enabled network meets this limitation by, *e.g.*, establishing a "Route64 TLV" link between the association unit and the station according to the Thread Specification following receipt of the Child ID Response.

114. Limitation F of Claim 1 requires "the operating route associated with the station running through a source of the association grant in response to said one of the association request messages." Use of the '717 Accused Instrumentalities in a Thread-enabled network meets this limitation. For example, the Response TLV runs through the Parent Router.

115. Limitation G of Claim 1 requires "if the source of the association grant is not the association unit, the operating route between the source of the association grant and the association unit previously established for the source of the association grant in response to association request messages and/or the association grants for intermediate stations." Use of the '717 Accused Instrumentalities in a Thread-enabled wireless network meets this limitation. For example, access to the network may be provided by a device other than the Parent, in which case the Thread Device will still be able to access the Parent via the operating route.

116. Limitation H of Claim 1 requires that "the association unit is configured to transmit association grants as responses to received ones of the association request messages." A Thread-enabled wireless network including the '717 Accused Instrumentalities meets this limitation because, *e.g.*, a router in such a network is configured to allocate Child IDs to new '717 Accused Instrumentalities in the network.

117. Limitation I of Claim 1 requires that "at least part of the stations is configured to become active to transmit association grants as responses to received ones of the association request messages, but only after switching to the associated state upon an association grant from the association unit or one of the at least part of the stations switched to the associated state directly

or indirectly based on an association grant from the association unit." The '717 Accused Instrumentalities meet this limitation because, *e.g.*, they are configured to act as routers within a Thread-enabled wireless network.

118. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '717 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of the '717 Accused Instrumentalities to directly infringe one or more claims of the '717 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '717 Patent and its infringement thereof no later than the filing of this Complaint, and (b) Amazon intentionally causes, urges, or encourages users of the '717 Accused Instrumentalities to take action that, when taken, directly infringe one or more claims of the '717 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use the '717 Accused Instrumentalities, including infringing uses thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of the '717 Accused Instrumentalities to directly infringe one or more claims of the '717 Patent, and users thereof directly infringe one or more claims of the '717 Patent. For instance, at a minimum, Amazon has supplied and continues to supply the '717 Accused Instrumentalities to customers while knowing that use thereof will infringe one or more claims of the '717 Patent.

119. Due to Amazon's knowledge of the '717 Patent and of its infringement thereof no later than the filing of this Complaint, Amazon's infringement of the '717 Patent has been and continues to be willful since no later than that date.

120. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

121. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

COUNT III
(Patent Infringement under 35 U.S.C. § 271 of the '120 Patent)

122. K.Mizra incorporates paragraphs 1 through 121 as though fully set forth herein.

123. The '120 Patent includes 16 claims.

124. Amazon has directly infringed one or more claims of the '120 Patent by making, importing, using, offering for sale, and/or selling Amazon SageMaker AI, all in violation of 35 U.S.C. § 271(a).

125. Based on publicly available information, Amazon SageMaker AI satisfies every element of at least Claim 1 of the '120 Patent.

126. Claim 1 of the '120 Patent states:

> [preamble] A method of determining hyperparameters of a classifier in a machine learning system by iteratively producing an estimate of a target hyperparameter vector, each iteration comprising the steps of::
>
> [A] drawing a random sample of hyperparameter vectors from a set of possible hyperparameter vectors,
>
> [B] updating the estimate of the target hyperparameter vector by using the random sample, and
>
> [C] selecting, from the random sample of hyperparameter vectors, a hyperparameter vector producing a best result in the present and any previous iterations, and
>
> [D] wherein the step of updating the estimate of the target hyperparameter vector uses said hyperparameter vector producing the best result.

(Ex. 3 at 7:14-28.)

127. As for the preamble of Claim 1, to the extent that it is determined to be limiting, Amazon SageMaker AI provides the features described in the preamble, which recites a "method of determining hyperparameters of a classifier in a machine learning system by iteratively producing an estimate of a target hyperparameter vector." Amazon SageMaker AI performs a "hyperparameter tuning" function called "automatic model tuning (AMT)."

> Amazon SageMaker AI automatic model tuning (AMT) finds the best version of a model by running many training jobs on your dataset. Amazon SageMaker AI automatic model tuning (AMT) is also known as hyperparameter tuning. To do this, AMT uses the algorithm and ranges of hyperparameters that you specify. It then chooses the hyperparameter values that creates a model that performs the best, as measured by a metric that you choose.

("Automatic model tuning with SageMaker AI, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning.html (last accessed January 2026 and incorporated by reference).) AMT "finds the best version of a model by running many training jobs on [a] dataset," "launch[ing] training jobs to find an optimal configuration of hyperparameters." (*Id.*)

> A hyperparameter is a high-level parameter that influences the learning process during model training. To get the best model predictions, you can optimize a hyperparameter configuration or set hyperparameter values. The process of finding an optimal configuration is called hyperparameter tuning. To configure and launch a hyperparameter tuning job, complete the steps in these guides.

> The hyperparameter tuning job will launch training jobs to find an optimal configuration of hyperparameters. These training jobs should be configured using the SageMaker AI `CreateHyperParameterTuningJob` API.

(Configure and Launch a Hyperparameter Tuning Job, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning-ex-tuning-job.html

(last accessed January 2026 and incorporated by reference).) Accordingly, and to the extent that the preamble of Claim 1 is somehow limiting, Amazon SageMaker AI would meet the limitation.

128. Limitation A of Claim 1 requires "drawing a random sample of hyperparameter vectors from a set of possible hyperparameter vectors." Amazon SageMaker AI meets this limitation. For example, Amazon SageMaker AI includes a "random search" function in which "hyperparameter tuning chooses a random combination of hyperparameter values in the ranges that you specify for each training job it launches."

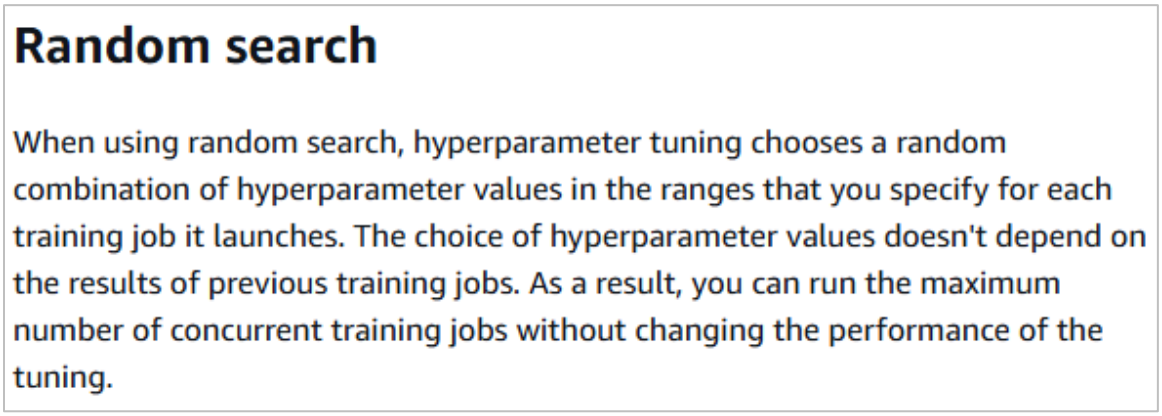

Random search

When using random search, hyperparameter tuning chooses a random combination of hyperparameter values in the ranges that you specify for each training job it launches. The choice of hyperparameter values doesn't depend on the results of previous training jobs. As a result, you can run the maximum number of concurrent training jobs without changing the performance of the tuning.

(Understand the hyperparameter tuning strategies available in Amazon SageMaker AI, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning-how-it-works.html (last accessed January 2026 and incorporated by reference).)

129. Limitation B of Claim 1 requires "updating the estimate of the target hyperparameter vector by using the random sample," and limitation C of Claim 1 requires "selecting, from the random sample of hyperparameter vectors, a hyperparameter vector producing a best result in the present and any previous iterations." Amazon SageMaker AI meets these limitations. For example, Amazon SageMaker AI "chooses the hyperparameter values that result in a model that performs the best," and "[w]hile the hyperparameter tuning job is in progress, the best training job is the one that has returned the best objective metric so far."

> The following section explains how to use an algorithm resource to run a hyperparameter tuning job in Amazon SageMaker AI. A hyperparameter tuning job finds the best version of a model by running many training jobs on your dataset using the algorithm and ranges of hyperparameters that you specify. It then chooses the hyperparameter values that result in a model that performs the best, as measured by a metric that you choose. For more information, see Automatic model tuning with SageMaker AI.

(Use an Algorithm to Run a Hyperparameter Tuning Job, available at https://docs.aws.amazon.com/sagemaker/latest/dg/sagemaker-mkt-algo-tune.html (last accessed January 2026 and incorporated by reference).)

> A hyperparameter tuning job uses the objective metric that each training job returns to evaluate training jobs. While the hyperparameter tuning job is in progress, the best training job is the one that has returned the best objective metric so far. After the hyperparameter tuning job is complete, the best training job is the one that returned the best objective metric.

(Configure and Launch a Hyperparameter Tuning Job, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning-ex-tuning-job.html (last accessed January 2026 and incorporated by reference).) This suggests that, in the "random search" tuning method, the optimal hyperparameter combination is chosen from the random sample in each training job/iteration. Amazon further explains that "[u]sing information from previous hyperparameter tuning jobs can help increase the performance of the new hyperparameter tuning job by making the search for the best combination of hyperparameters more efficient," further indicating that Amazon SageMaker AI selects the best hyperparameter combinations for each training job based on the previous iteration.

> Use warm start to start a hyperparameter tuning job using one or more previous tuning jobs as a starting point. The results of previous tuning jobs are used to inform which combinations of hyperparameters to search over in the new tuning job. Hyperparameter tuning uses either Bayesian or random search to choose combinations of hyperparameter values from ranges that you specify. For more information, see Understand the hyperparameter tuning strategies available in Amazon SageMaker AI. Using information from previous hyperparameter tuning jobs can help increase the performance of the new hyperparameter tuning job by making the search for the best combination of hyperparameters more efficient.

(Run a Warm Start Hyperparameter Tuning Job, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning-warm-start.html (last accessed January 2026 and incorporated by reference).)

130. Limitation D of Claim 1 requires that " the step of updating the estimate of the target hyperparameter vector uses said hyperparameter vector producing the best result." Amazon SageMaker AI meets this limitation, as the best training job hyperparameters are stored separately, indicating that the best training job hyperparameters ("the estimate of the target hyperparameter vector") are updated whenever a new optimal set is found.

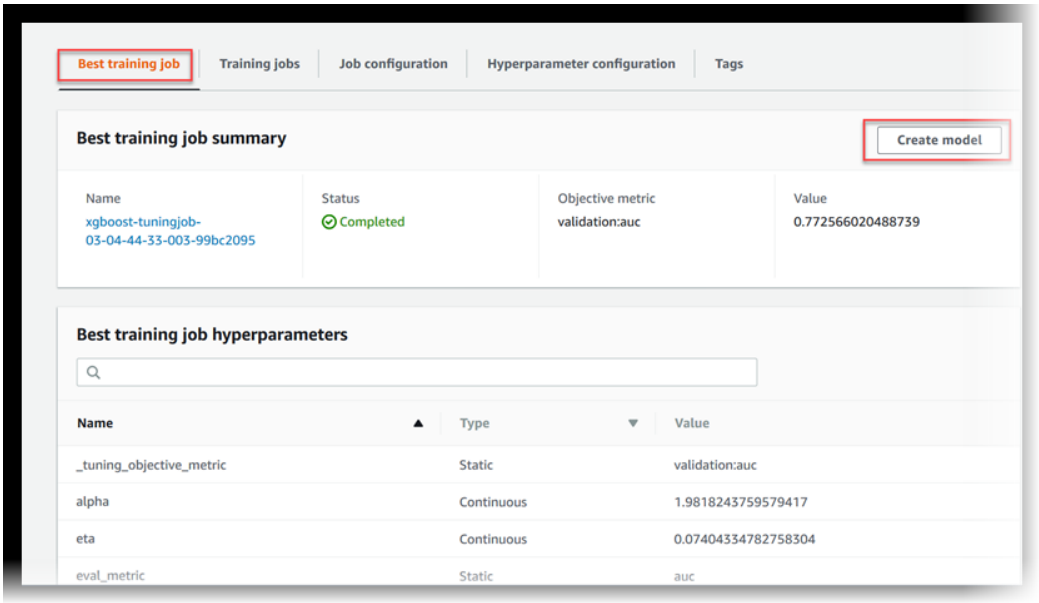

(Configure and Launch a Hyperparameter Tuning Job, available at https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning-ex-tuning-job.html (last accessed January 2026 and incorporated by reference).)

131. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '120 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of Amazon SageMaker AI to directly infringe one or more claims of the '120 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '120 Patent and its infringement thereof no later than the filing of this Complaint, and (b) Amazon intentionally causes, urges, or encourages users of Amazon SageMaker AI to take action that, when taken, directly infringe one or more claims of the '120 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use Amazon SageMaker AI, including infringing uses thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of Amazon SageMaker AI to directly infringe one or more claims of the '120 Patent, and users thereof directly infringe one or more claims of the '120 Patent. For instance, at a minimum, Amazon has supplied and continues to supply Amazon SageMaker AI to customers while knowing that use thereof will infringe one or more claims of the '120 Patent.

132. Due to Amazon's knowledge of the '120 Patent and of its infringement thereof no later than the filing of this Complaint, Amazon's infringement of the '120 Patent has been and continues to be willful since no later than that date.

133. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

134. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

COUNT IV
(Patent Infringement under 35 U.S.C. § 271 of the '259 Patent)

135. K.Mizra incorporates paragraphs 1 through 134 as though fully set forth herein.

136. The '259 Patent includes 23 claims.

137. Amazon has directly infringed one or more claims of the '259 Patent by making, importing, using, offering for sale, and/or selling Alexa Emergency Assist, all in violation of 35 U.S.C. § 271(a).

138. Based on publicly available information, Alexa Emergency Assist satisfies every element of at least Claim 1 of the '259 Patent.

139. Claim 1 of the '259 Patent states:

> [preamble] Method for detecting ticks on a device with which an audio signal is received in a noisy environment, from the audio signal which includes a stream of audio samples (y(n), (y′(n)), the method comprising the steps of:
>
> [A] applying the audio signal to a coarse tick detection processor arranged to decide whether the audio signal likely includes a tick; and
>
> [B] only if the course tick processor decides that the audio signal likely includes a tick, enabling a fine tick detection processor to decide, by more thorough processing of said audio signal, whether the audio signal represents a tick on the device,
>
> [C] wherein the fine tick detection is performed using a comparison with a previously trained set of reference properties that are characteristic for the device with which the audio signal is received.

(Ex. 4 at 7:14-28.)

140. As for the preamble of Claim 1, to the extent that it is determined to be limiting, Alexa Emergency Assist provides the features described in the preamble, which recites a "[m]ethod for detecting ticks on a device with which an audio signal is received in a noisy environment, from the audio signal which includes a stream of audio samples (y(n), (y′(n))." For

example, Alexa Emergency Assist can detect the sound of smoke alarms, carbon monoxide alarms, and glass breaking.

> **1. What sounds can Alexa detect?**
>
> Alexa Emergency Assist can detect and notify you about the sound of smoke alarms and carbon monoxide alarms, and glass breaking.

(Alexa Emergency Assist – Sound Detection FAQs, available at https://www.amazon.com/gp/help/customer/display.html?nodeId=TZhqo3nwWAuKfX0BBT (last accessed January 2026 and incorporated by reference).) Further, the previous version of Alexa Emergency Assist, Alexa Guard, "allow[ed] customers to adjust the sensitivity to accommodate the noise in their home environments." (The science behind an Amazon Echo feature that helped save a puppy, available at https://www.amazon.science/latest-news/the-science-behind-an-amazon-echo-feature-that-helped-save-a-puppy (last accessed January 2026 and incorporated by reference).) According to Amazon staff member "Roger B.," Alexa Emergency Assist includes "the top features . . . from Alexa Guard" in an improved form.

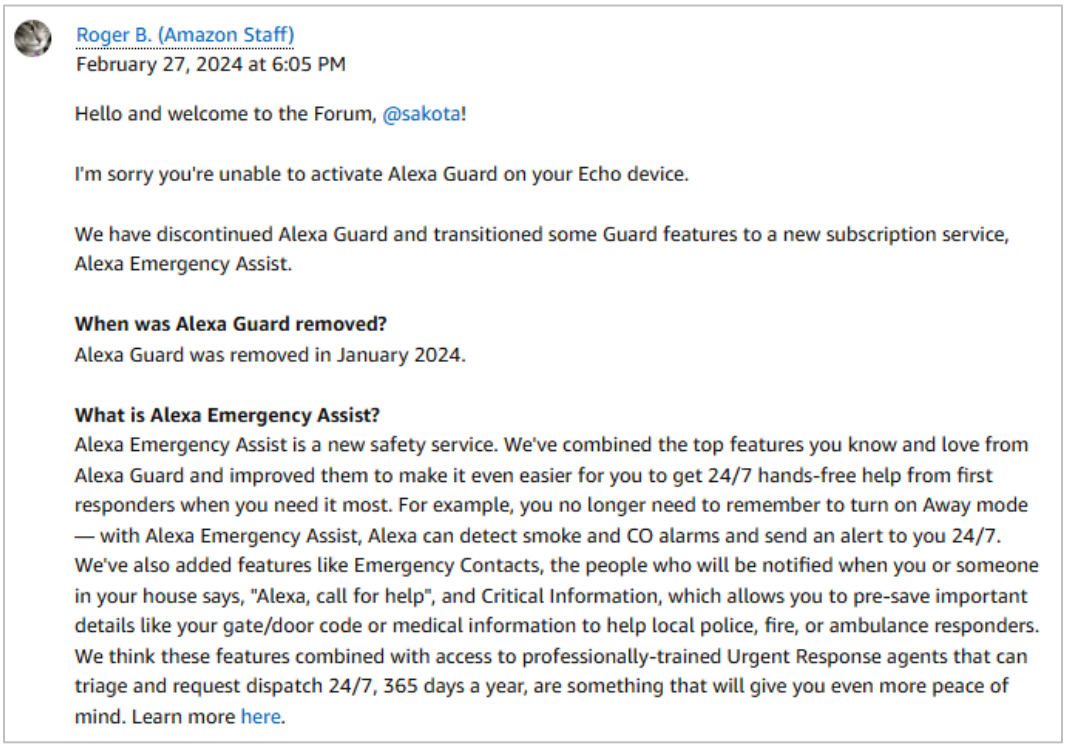

Roger B. (Amazon Staff)
February 27, 2024 at 6:05 PM

Hello and welcome to the Forum, @sakota!

I'm sorry you're unable to activate Alexa Guard on your Echo device.

We have discontinued Alexa Guard and transitioned some Guard features to a new subscription service, Alexa Emergency Assist.

When was Alexa Guard removed?
Alexa Guard was removed in January 2024.

What is Alexa Emergency Assist?
Alexa Emergency Assist is a new safety service. We've combined the top features you know and love from Alexa Guard and improved them to make it even easier for you to get 24/7 hands-free help from first responders when you need it most. For example, you no longer need to remember to turn on Away mode — with Alexa Emergency Assist, Alexa can detect smoke and CO alarms and send an alert to you 24/7. We've also added features like Emergency Contacts, the people who will be notified when you or someone in your house says, "Alexa, call for help", and Critical Information, which allows you to pre-save important details like your gate/door code or medical information to help local police, fire, or ambulance responders. We think these features combined with access to professionally-trained Urgent Response agents that can triage and request dispatch 24/7, 365 days a year, are something that will give you even more peace of mind. Learn more here.

(Alexa Guard, available at https://amazonforum.my.site.com/s/question/0D56Q0000Cw9TunSQE/alexa-guard (last accessed January 2026 and incorporated by reference).) Thus, on information and belief, Alexa Emergency Assist includes the environmental noise features of Alexa Guard.

141. Limitation A of Claim 1 requires "applying the audio signal to a coarse tick detection processor arranged to decide whether the audio signal likely includes a tick." Alexa Emergency Assist meets this limitation. The Guard/Emergency Assist technology "relies on two models applied in a two-step system, one on the device, another in the cloud." (The science behind an Amazon Echo feature that helped save a puppy, available at https://www.amazon.science/latest-news/the-science-behind-an-amazon-echo-feature-that-helped-save-a-puppy (last accessed January 2026 and incorporated by reference).) "[T]he majority of the audio data is processed and discarded by the neural network on the device," indicating that the first-step device-side detection involves "a coarse tick detection processor arranged to decide whether the audio signal likely includes a tick." (*Id.*)

142. Limitation B of Claim 1 requires, "only if the course tick processor decides that the audio signal likely includes a tick, enabling a fine tick detection processor to decide, by more thorough processing of said audio signal, whether the audio signal represents a tick on the device." Alexa Emergency Assist meets this limitation. For example, in the Guard/Emergency Assist technology, "audio is sent to the cloud only after running it through a device-side model." (*Id.*) Further, as noted above, "the majority of the audio data is processed and discarded by the neural network on the device," and "[o]nly potential triggers make it to the cloud" for the cloud-side processing. (*Id.*) The cloud-side processing (fine tick detection processor) conducts a "second verification step to confirm the on-device detection," using "a much more powerful recognition

system to filter out false triggers that might be linked to ambient noise around the home." (*Id.*) Thus, the Guard/Emergency Assist technology includes enabling a fine tick detection processor to decide, by more thorough processing of said audio signal, whether the audio signal represents a tick on the device.

143. Limitation C of Claim 1 requires that "the fine tick detection is performed using a comparison with a previously trained set of reference properties that are characteristic for the device with which the audio signal is received." On information and belief, the Guard/Emergency Assist cloud-side verification process involves a previously trained set of reference properties that are characteristic for the device with which the audio signal is received. For example, the device-side detection and processing "utilizes a recurrent neural network — a type of deep learning model that uses sequential data or time series data to learn — on the Echo device itself. The on-device detection works by converting the audio input into features that feed into a recurrent neural network (RNN)." (*Id.*) Further, Amazon "rented a warehouse and contracted a construction crew to break hundreds of windows . . . to build an authentic data set" to train the Guard/Emergency Assist technology. (*Id.*) Accordingly, on information and belief, the cloud-side verification process is performed using a comparison with a previously trained set of reference properties that are characteristic for the device with which the audio signal is received.

144. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '259 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of Alexa Emergency Assist to directly infringe one or more claims of the '259 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '259 Patent and its infringement thereof no later than the filing of this Complaint, and (b) Amazon intentionally causes, urges, or encourages users of Alexa Emergency Assist to

take action that, when taken, directly infringe one or more claims of the '259 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use Alexa Emergency Assist, including infringing uses thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of Alexa Emergency Assist to directly infringe one or more claims of the '259 Patent, and users thereof directly infringe one or more claims of the '259 Patent. For instance, at a minimum, Amazon has supplied and continues to supply Alexa Emergency Assist to customers while knowing that use thereof will infringe one or more claims of the '259 Patent.

145. Due to Amazon's knowledge of the '259 Patent and of its infringement thereof no later than the filing of this Complaint, Amazon's infringement of the '259 Patent has been and continues to be willful since no later than that date.

146. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

147. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

COUNT V
(Patent Infringement under 35 U.S.C. § 271 of the '649 Patent)

148. K.Mizra incorporates paragraphs 1 through 147 as though fully set forth herein.

149. The '649 Patent includes 26 claims.

150. Amazon has directly infringed one or more claims of the '649 Patent by making, importing, using, offering for sale, and/or selling Amazon Echo devices incorporating Echo Spatial Perception, all in violation of 35 U.S.C. § 271(a).

151. Based on publicly available information, Amazon Echo devices incorporating Echo Spatial Perception satisfy every element of at least Claim 1 of the '649 Patent.

152. Claim 1 of the '649 Patent states:

> [preamble] A method of identifying proximate devices, the method comprising the steps of:
>
> [A] recording, in each device, an audio signal so as to produce respective audio samples,
>
> [B] detecting, in each device, a detected sensory identifier, and
>
> [C] transmitting the respective audio samples from one or more of the devices to a server or another of the devices, wherein the server or device receiving the respective audio samples comprises a processor, a memory and software configured to perform the remaining steps of
>
> [D] using said detected sensory identifier for matching the devices, using said detecting as a trigger to select a common time interval of audio samples at a predetermined relative temporal position offset with respect to a time point of the detected sensory identifier;
>
> [E] comparing respective audio samples recorded during the common time interval by different ones of the devices with each other, and
>
> [F] deciding that the devices are proximate only if the respective audio samples are sufficiently similar.

(Ex. 5 at 9:18-39.)

153. As for the preamble of Claim 1, to the extent that it is determined to be limiting, Amazon Echo devices incorporating Echo Spatial Perception provides the features described in the preamble, which recites a "method of identifying proximate devices." Echo Spatial Perception

"enable[s] devices to collaboratively determine which is closes to the user" and thus performs a "method of identifying proximate devices."

> Amazon has officially moved ESP to the cloud, which should ensure that all devices, Amazon or non-Amazon, can use it. Users don't even have to download any software updates to begin using the feature; it'll come on board automatically. Literally any Alexa device able to connect to the internet can now start using the feature. ESP will enable these devices to collaboratively determine which is closest to the user. Amazon was short on specifics, but presumably this is done by basing which noise-canceling array detects the strongest audio signal.

(How Amazon Ensures Only The Closest Alexa Device Responds To Your Voice Commands, available at https://www.techtimes.com/articles/232691/20180726/how-amazon-ensures-only-the-closest-alexa-device-responds-to-your-voice-commands.htm (last accessed January 2026 and incorporated by reference).)

154. Limitation A of Claim 1 requires "recording, in each device, an audio signal so as to produce respective audio samples." The Echo Spatial Perception process meets this limitation because any given Echo device in proximity will record audio input following detection of the chosen wake word.

> **Fact: By default, Echo devices are designed to record audio only after the device detects your chosen wake word (or if the action button is pushed).**
>
> Until your Echo device detects your chosen wake word—current options include “Alexa,” “Amazon,” “Echo,” “Ziggy,” and “Computer”—Alexa does not record or store what you or any other person in the room has said. This is made possible by built-in technology called “keyword spotting” that matches spoken audio to the wake word’s acoustic patterns. Echo devices are designed to listen for the sound waves of the wake word while ignoring everything else. This is why your device can detect the wake word without listening in on your personal conversations.

(Busting the four biggest myths about Amazon's Alexa, available at https://www.aboutamazon.eu/busting-the-four-biggest-myths-about-amazons-alexa (last accessed January 2026 and incorporated by reference).)

155. Limitation B of Claim 1 requires "detecting, in each device, a detected sensory identifier." The Echo Spatial Perception process meets this limitation because each Echo device will detect a "wake word" (sensory identifier). (*Id.*)

156. Limitation C of Claim 1 requires "transmitting the respective audio samples from one or more of the devices to a server or another of the devices, wherein the server or device receiving the respective audio samples comprises a processor, a memory and software." The Echo Spatial Perception process meets this limitation because spoke audio inputs are transmitted to "Amazon's secure cloud" by Echo devices.

> When you speak to Alexa, your request is sent to Amazon's secure cloud so that Alexa can respond to you. Customers are in control of their voice recordings. You can see exactly what Alexa heard and sent to Amazon's cloud in the Alexa app by visiting Settings > Alexa Privacy > Review Voice History. Or you can look online at your Alexa privacy settings. Listen to the actual audio recorded or manage your voice history by choosing to delete specific recordings—sorting by date, device, or profile. You can delete them all at once, have them deleted on an ongoing basis, or not have them saved at all.

(*Id.*) On information and belief, Amazon's cloud servers comprise a processor, a memory and software configured to perform the remaining steps described below. In particular, Echo Spatial Perception is implemented on the cloud, and thus Amazon's cloud servers perform the Echo Spatial Perception process.

> *There are millions of customers with more than one Alexa-enabled device in their home, and for these customers one of the most important features we've built is Echo Spatial Perception (ESP). The ESP feature provides a frustration-free experience by identifying the best device to respond when interacting with Alexa in a multi-device environment. The technology determines which device is closest to the customer, and helps Alexa respond through that device.... Today we're announcing that ESP has moved to the cloud, introducing several benefits to developers. With Cloud ESP, all existing and new Alexa-enabled devices get the ESP feature with no requirement for device-side software changes. Cloud ESP also offers superior accuracy – even in noisy environments – and because Alexa is always getting smarter, the feature will continue to improve over time.*

(Echo Spatial Perception moves to the cloud, available at https://vowe.net/archives/017276.html (last accessed January 2026 and incorporated by reference).)

157. Limitation D of Claim 1 requires "using said detected sensory identifier for matching the devices, using said detecting as a trigger to select a common time interval of audio samples at a predetermined relative temporal position offset with respect to a time point of the detected sensory identifier." Amazon's cloud servers performing the Echo Spatial Perception process perform this step by, for example, using the volume and proximity of spoken requests to determine which Echo device is closest to the user.

> I'd like to inform you that, your firmware version is updated to recent one, which requires for ESP (Echo Spatial Perception) (i.e. 4008). With ESP, If you have more than one Alexa device using the same wake word (for example, two Echo devices using "Alexa" as a wake word), Alexa chooses the device closest to you to respond to your request using ESP (Echo Spatial Perception). ESP uses the volume and proximity of the spoken request. Nearby Alexa devices that happen to recognize the same wake word won't respond to the request.

(Amazon Echo rolls out Echo Spatial Perception, available at https://community.smartthings.com/t/amazon-echo-rolls-out-echo-spatial-perception/59873/34?page=2 (last accessed January 2026 and incorporated by reference).)

> Amazon has officially moved ESP to the cloud, which should ensure that all devices, Amazon or non-Amazon, can use it. Users don't even have to download any software updates to begin using the feature; it'll come on board automatically. Literally any Alexa device able to connect to the internet can now start using the feature. ESP will enable these devices to collaboratively determine which is closest to the user. Amazon was short on specifics, but presumably this is done by basing which noise-canceling array detects the strongest audio signal.

(How Amazon Ensures Only The Closest Alexa Device Responds To Your Voice Commands, available at https://www.techtimes.com/articles/232691/20180726/how-amazon-ensures-only-the-closest-alexa-device-responds-to-your-voice-commands.htm (last accessed January 2026 and incorporated by reference).) This indicates that all Echo devices that detect the wake word align their audio processing to the same timeframe (*e.g.*, the period of the user's query following the wake word) so that their recordings can be compared on an equal basis. In other words, the wake word detection triggers the process of evaluating that request audio across devices. Accordingly, the Echo Spatial Perception process meets limitation D of claim 1.

158. Limitation E of Claim 1 requires "comparing respective audio samples recorded during the common time interval by different ones of the devices with each other." Amazon's cloud servers performing the Echo Spatial Perception process perform this step. As discussed above with respect to limitation D, the Echo Spatial Perception process involves a comparison of audio samples based on, *e.g.*, volume to determine the closest device. Further, "Cloud ESP offers superior accuracy – even in noisy environments," indicating that the system performs a more complex comparison and analyzes audio characteristics beyond volume.

> *There are millions of customers with more than one Alexa-enabled device in their home, and for these customers one of the most important features we've built is Echo Spatial Perception (ESP). The ESP feature provides a frustration-free experience by identifying the best device to respond when interacting with Alexa in a multi-device environment. The technology determines which device is closest to the customer, and helps Alexa respond through that device.... Today we're announcing that ESP has moved to the cloud, introducing several benefits to developers. With Cloud ESP, all existing and new Alexa-enabled devices get the ESP feature with no requirement for device-side software changes. Cloud ESP also offers superior accuracy – even in noisy environments – and because Alexa is always getting smarter, the feature will continue to improve over time.*

(Echo Spatial Perception moves to the cloud, available at https://vowe.net/archives/017276.html (last accessed January 2026 and incorporated by reference).)

159. Limitation F of Claim 1 requires "deciding that the devices are proximate only if the respective audio samples are sufficiently similar." An Amazon cloud server performing the Echo Spatial Perception process meets this limitation because the process identifies devices to be distinguished based on proximity only if the devices respond to the same audio input. In other words, Echo Spatial Perception decides that two devices are proximate only if they detect/record the same audio samples.

160. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '649 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of Amazon Echo devices incorporating Echo Spatial Perception to directly infringe one or more claims of the '649 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '649 Patent and its infringement thereof no later than the filing of this Complaint, and (b) Amazon intentionally causes, urges, or encourages users of Amazon Echo devices incorporating Echo Spatial Perception to take action that, when taken, directly infringe one or more claims of the '649 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use Amazon Echo devices incorporating Echo Spatial Perception, including infringing uses

thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of Amazon Echo devices incorporating Echo Spatial Perception to directly infringe one or more claims of the '649 Patent, and users thereof directly infringe one or more claims of the '649 Patent. For instance, at a minimum, Amazon has supplied and continues to supply Amazon Echo devices incorporating Echo Spatial Perception to customers while knowing that use thereof will infringe one or more claims of the '649 Patent.

161. Due to Amazon's knowledge of the '649 Patent and of its infringement thereof no later than the filing of this Complaint, Amazon's infringement of the '649 Patent has been and continues to be willful since no later than that date.

162. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

163. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

COUNT VI
(Patent Infringement under 35 U.S.C. § 271 of the '282 Patent)

164. K.Mizra incorporates paragraphs 1 through 163 as though fully set forth herein.

165. The '282 Patent includes 22 claims.

166. Amazon has directly infringed one or more claims of the '282 Patent by making, importing, using, offering for sale, and/or selling Amazon CloudWatch, all in violation of 35 U.S.C. § 271(a).

167. Based on publicly available information, Amazon CloudWatch satisfies every element of at least Claim 1 of the '282 Patent.

168. Claim 1 of the '282 Patent states:

> [preamble] A method, comprising:
>
> [A] receiving, at a first application server instance selected from a plurality of application server instances based on a load balancing process, first adapter processed information from a first adapter, wherein the first adapter processed information comprises event information received by the first adapter from a network element and processed by the first adapter based on a first communication protocol;
>
> [B] processing, by the first application server instance, the first adapter processed information based on an event management service to produce application processed information;
>
> [C] sending, by the first application server instance, the application processed information to a gateway device, wherein the gateway device is one of a plurality of gateway devices respectively associated with the plurality of application server instances and is configured to transfer the application processed information to a second adapter of a plurality of second adapters configured to process the application processed information based on a second communication protocol to produce second adapter processed information and transfer the second adapter processed information to an operation support system device; and
>
> [D] in response to determining that the first application server instance has become disabled, facilitating establishing an association between the first adapter and a second application server instance of the plurality of application server instances and between the gateway device and the second application server instance.

(Ex. 6 at 10:59-11:22.)

169. As for the preamble of Claim 1, to the extent that it is determined to be limiting, Amazon CloudWatch provides the features described in the preamble because Amazon CloudWatch performs a "method."

170. Limitation A of Claim 1 requires "receiving, at a first application server instance selected from a plurality of application server instances based on a load balancing process, first adapter processed information from a first adapter, wherein the first adapter processed information comprises event information received by the first adapter from a network element and processed by the first adapter based on a first communication protocol." Amazon CloudWatch performs this step. The Amazon CloudWatch agent can be installed on "Amazon EC2 instances" that will then be capable of "Elastic Load Balancing," which enables the launch of multiple EC2 instances and distribution of traffic between them.

> You can collect metrics from servers by installing the CloudWatch agent on the server. You can install the agent on both Amazon EC2 instances and on-premises servers. You can also install the agent on computers running Linux, Windows Server, or macOS. If you install the agent on an Amazon EC2 instance, the metrics the agent collects are in addition to the metrics enabled by default on Amazon EC2 instances. For information about installing the CloudWatch agent on an instance, see Collect metrics, logs, and traces using the CloudWatch agent. You can use this section to learn about metrics the CloudWatch agent collects.

(Metrics collected by the CloudWatch agent, available at https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/metrics-collected-by-CloudWatch-agent.html (last accessed January 2026 and incorporated by reference).)

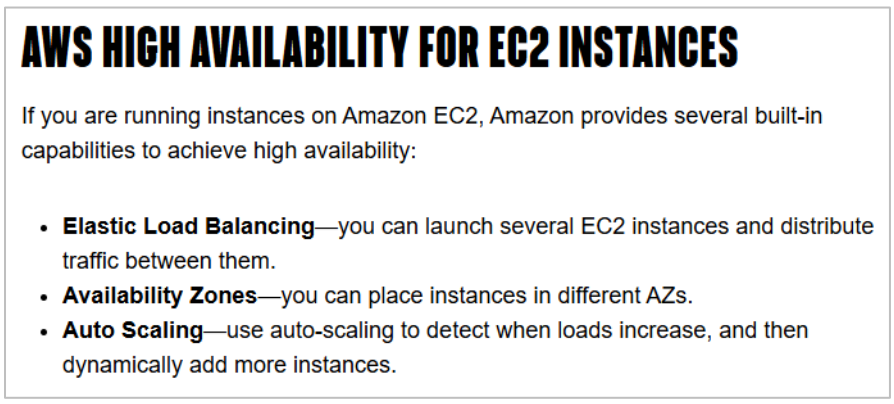

(Understanding AWS High Availability: Compute, SQL and Storage, available at https://www.netapp.com/learn/understanding-aws-high-availability-compute-sql-and-storage/ (last accessed January 2026 and incorporated by reference).)

171. CloudWatch is integrated with an SNMP agent such Logstash ("first adapter"), which collects information from the monitored devices such as routers ("network elements"), and, based on a communication protocol such as SNMP, converts it into metric information concerning monitored devices ("first adapter processed information").

> In this solution, Logstash acts as a relay server that sits between the managed devices that you want to monitor and Amazon CloudWatch. I use two Logstash plugins: SNMP (input plugin) and CloudWatch (output plugin). The input plugin collects information (using SNMP polling) from the monitored devices. These devices must be running SNMP daemons, listening on standard port 161 (Poll). After Logstash processes the information, the output plugin forwards the metrics to CloudWatch using standard AWS APIs. I use AWS Secrets Manager and Logstash keystore for SNMP credential management. You can expand this solution by using Logstash Snmptrap input plugin.

(SNMP monitoring using Amazon CloudWatch and Elastic Logstash, available at https://aws.amazon.com/blogs/mt/snmp-monitoring-using-amazon-cloudwatch-and-elastic-logstash/ (last accessed January 2026 and incorporated by reference).)

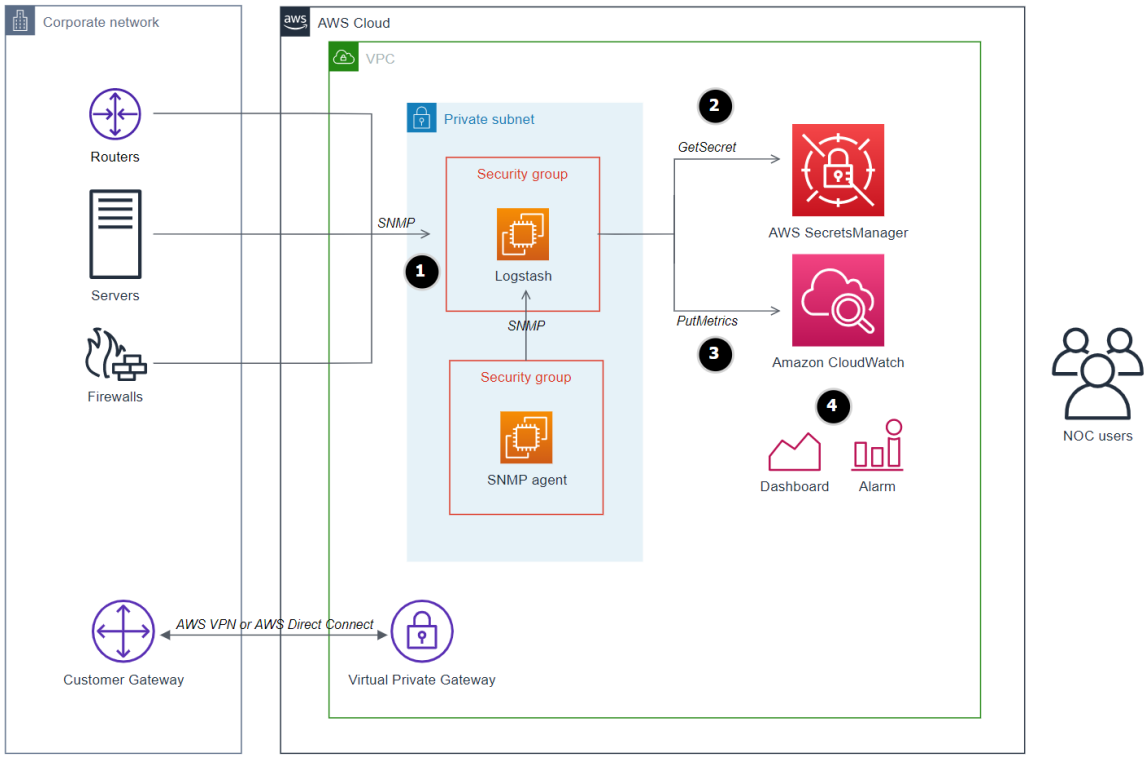

(*Id.*) CloudWatch then receives the metric information from Logstash.

172. Limitation B of Claim 1 requires "processing, by the first application server instance, the first adapter processed information based on an event management service to produce application processed information." CloudWatch performs this step. For example, CloudWatch creates ("produces") alarms/events ("application processed information") based on the metrics ("first adapter processed information").

> You can use an *alarm* to automatically initiate actions on your behalf. An alarm watches a single metric over a specified time period, and performs one or more specified actions, based on the value of the metric relative to a threshold over time. The action is a notification sent to an Amazon SNS topic or an Auto Scaling policy. You can also add alarms to dashboards.

(Amazon CloudWatch User Guide at p. 125, available at https://docs.aws.amazon.com/pdfs/AmazonCloudWatch/latest/monitoring/acw-ug.pdf (last accessed January 2026 and incorporated by reference).)

173. Limitation C of Claim 1 requires "sending, by the first application server instance, the application processed information to a gateway device, wherein the gateway device is one of a plurality of gateway devices respectively associated with the plurality of application server instances and is configured to transfer the application processed information to a second adapter of a plurality of second adapters configured to process the application processed information based on a second communication protocol to produce second adapter processed information and transfer the second adapter processed information to an operation support system device." CloudWatch performs this step. For example, the alarms/events referenced above with respect to limitation B are sent to Amazon Data Firehose ("gateway device") either directly or via the Amazon SNS.

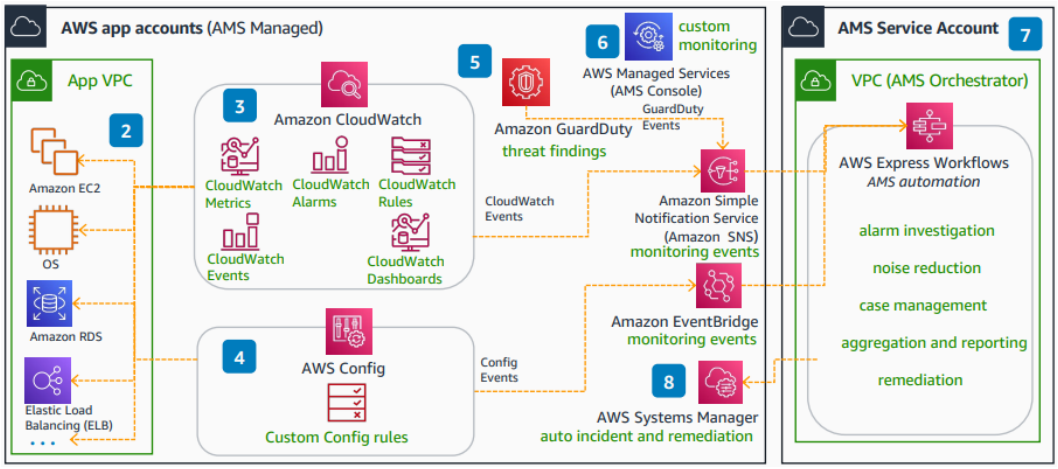

(AWS Managed Services (AMS) for Operational Excellence at p. 3, available at https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/AWS-managed-services-for-operational-excellence-ra.pdf (last accessed January 2026 and incorporated by reference).)

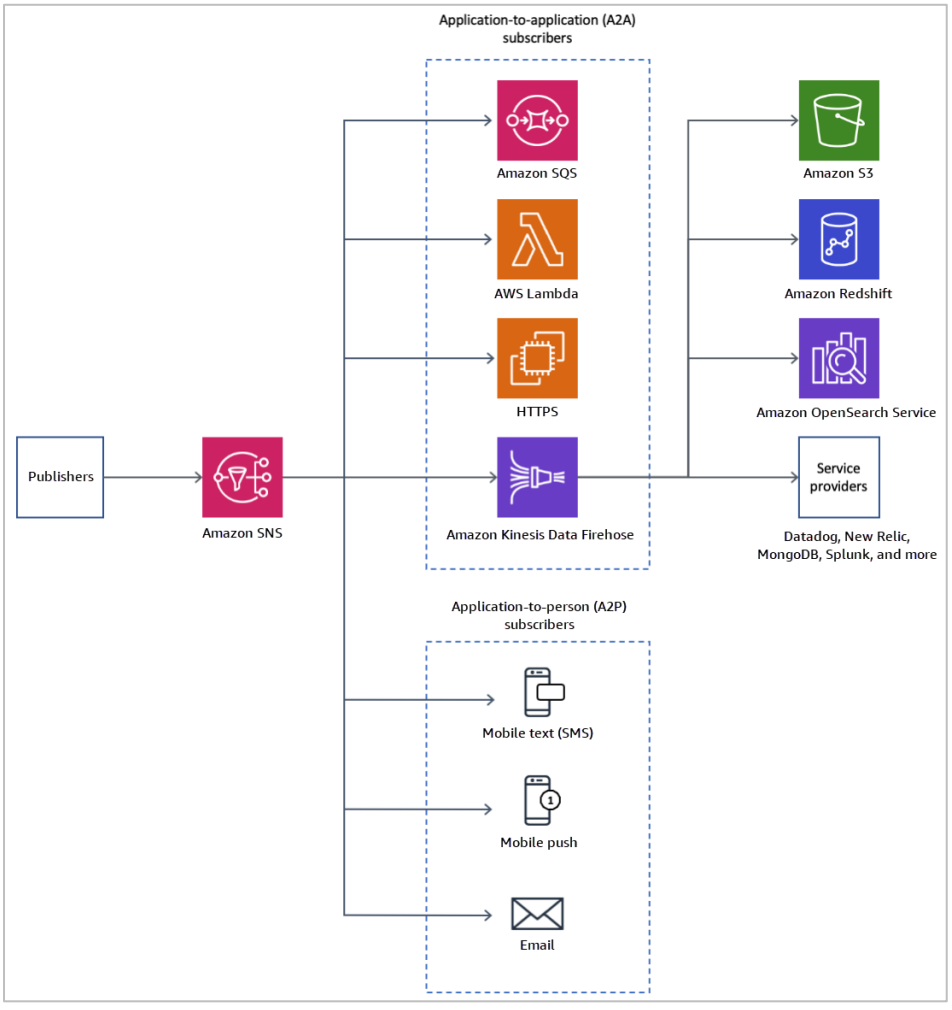

(What is Amazon SNS?, available at https://docs.aws.amazon.com/sns/latest/dg/welcome.html (last accessed January 2026 and incorporated by reference).) Amazon Data Firehose forwards the alarms/events to Amazon OpenSearch Service ("operation support system") using, e.g., interfaces ("second adapter") that support the HTTPS protocol ("second communication protocol").

> Amazon Data Firehose is a fully managed service for delivering real-time streaming data to destinations such as Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, Amazon OpenSearch Serverless, Splunk, Apache Iceberg Tables, and any custom HTTP endpoint or HTTP endpoints owned by supported third-party service providers, including Datadog, Dynatrace, LogicMonitor, MongoDB, New Relic, Coralogix, and Elastic. With Amazon Data Firehose, you don't need to write

(What is Amazon Data Firehose?, available at https://docs.aws.amazon.com/firehose/latest/dev/what-is-this-service.html (last accessed January 2026 and incorporated by reference).)

174. Limitation D of Claim 1 requires "in response to determining that the first application server instance has become disabled, facilitating establishing an association between the first adapter and a second application server instance of the plurality of application server instances and between the gateway device and the second application server instance." CloudWatch performs this step. For example, the AWS architecture implements Availability Zones, which support failover between zones, indicating that when one CloudWatch agent (deployed on an EC2 instance) goes down, another CloudWatch agent takes its place by, *e.g.*, establishing a connection between the new instance and the SNMP agents/Logstash and the Amazon Data Firehose ("gateway").

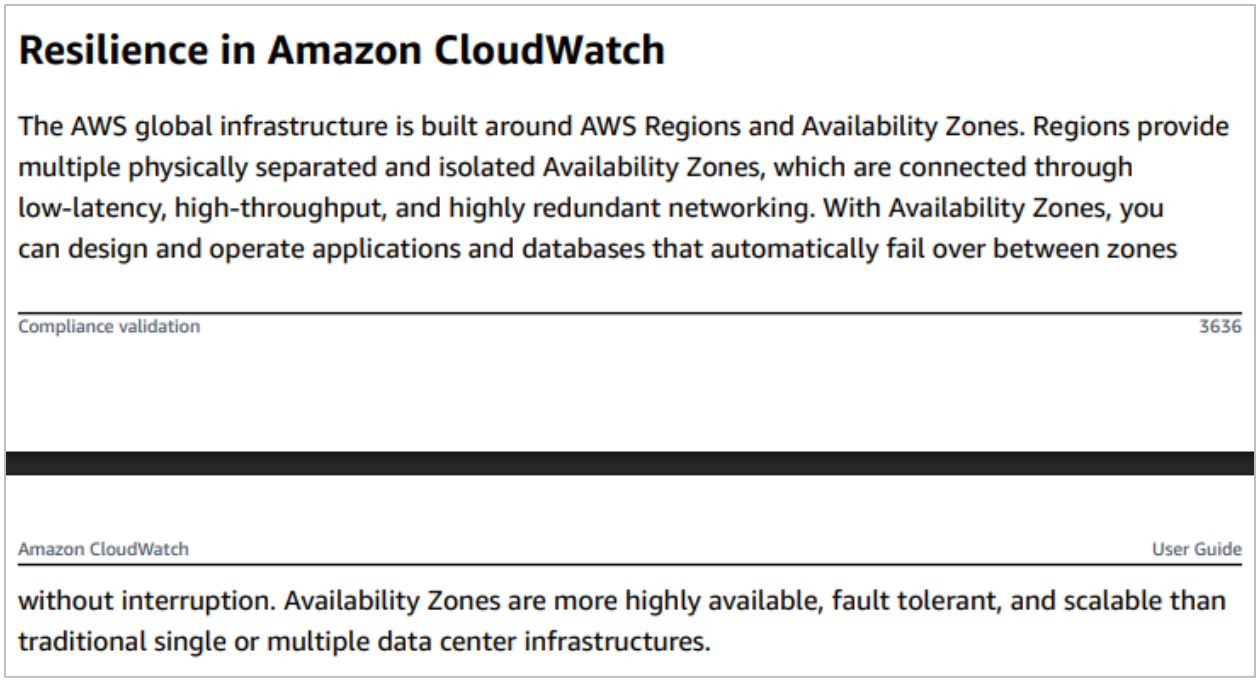

Resilience in Amazon CloudWatch

The AWS global infrastructure is built around AWS Regions and Availability Zones. Regions provide multiple physically separated and isolated Availability Zones, which are connected through low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones

without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

(Amazon CloudWatch User Guide at pp. 3636-37, available at https://docs.aws.amazon.com/pdfs/AmazonCloudWatch/latest/monitoring/acw-ug.pdf (last accessed January 2026 and incorporated by reference).)

175. Additionally, and/or alternatively, Amazon has indirectly infringed and continues to indirectly infringe one or more of the claims of the '282 Patent, in violation of 35 U.S.C. § 271(b) by actively inducing users of Amazon CloudWatch to directly infringe one or more claims of the '282 Patent. For example, (a) Amazon had actual knowledge of or was willfully blind to the existence of the '282 Patent and its infringement thereof no later than the filing of this Complaint, and (b) Amazon intentionally causes, urges, or encourages users of Amazon CloudWatch to take action that, when taken, directly infringe one or more claims of the '282 Patent. Amazon's encouragement is accomplished by promoting, advertising, and instructing customers and potential customers to use Amazon CloudWatch, including infringing uses thereof. Amazon knows (based on the claim chart previously provided by K.Mizra and/or after reading this Complaint should know) that its actions will induce users of Amazon CloudWatch to directly infringe one or more

claims of the '282 Patent, and users thereof directly infringe one or more claims of the '282 Patent. For instance, at a minimum, Amazon has supplied and continues to supply Amazon CloudWatch to customers while knowing that use thereof will infringe one or more claims of the '282 Patent.

176. Due to Amazon's knowledge of the '282 Patent and of its infringement thereof no later than the filing of this Complaint, Amazon's infringement of the '282 Patent has been and continues to be willful since no later than that date.

177. Amazon's acts of infringement have occurred within this District and elsewhere throughout the United States.

178. As a result of Amazon's infringing conduct, K.Mizra has suffered damages. Amazon is liable to K.Mizra in an amount that adequately compensates K.Mizra for Amazon's infringement in an amount that is no less than a fully paid-up, lump-sum, reasonable royalty, together with interest and costs as fixed by this Court under 25 U.S.C. § 284.

REQUEST FOR RELIEF

WHEREFORE, K.Mizra respectfully requests the Court find in its favor and against Amazon, and that the Court grant K.Mizra at least the following relief:

A. Judgment that Amazon has directly infringed, literally and/or under the DoE, one or more claims of the Asserted Patents;

B. Awarding damages to K.Mizra in an amount to be proven at trial and in the form of a fully paid-up, lump sum, reasonable royalty that takes into account and runs through expiration of the Asserted Patents;

C. Awarding enhanced damages, as appropriate, under 35 U.S.C. § 284;

D. Awarding K.Mizra's costs (including disbursements) and declaring this an exceptional case and awarding K.Mizra its attorneys' fees in accordance with 35 U.S.C. § 285;

E. Pre-judgment and post-judgment interest at the maximum rate permitted by law on the damages caused to K.Mizra by reason of Amazon's infringing activities and other conduct complaint of herein; and

F. Awarding such other and further relief as this Court deems just and proper under the circumstances.

DEMAND FOR JURY TRIAL

Pursuant to Fed. R. Civ. P. 38(b), K.Mizra hereby demands a trial by jury on all issues so triable.

Dated: February 10, 2026

Respectfully submitted,

By: */s/ Michael C. Smith*

Michael C. Smith
Texas Bar No. 18650410
michael.smith@solidcounsel.com
Scheef & Stone, LLP
113 E. Austin Street
Marshall, TX 75670
(903) 938-8900

Robert R. Brunelli*
CO State Bar No. 20070
rbrunelli@sheridanross.com
Brian S. Boerman*
CO State Bar No. 50834
bboerman@sheridanross.com
Matthew C. Holohan*
CO State Bar No. 40996
mholohan@sheridanross.com
SHERIDAN ROSS P.C.
1560 Broadway, Suite 1200
Denver, CO 80202
Telephone: 303-863-9700

Facsimile: 303-863-0223
litigation@sheridanross.com

Of Counsel:
Claire A. Henry
Texas Bar No. 24053063
Miller Fair Henry PLLC
1507 Bill Owens Parkway
Longview, TX 75604
Telephone: (903) 757-6400
Facsimile: (903) 757-2323
claire@millerfairhenry.com

**To be admitted pro hac vice*
Attorneys for Plaintiff K.Mizra LLC