UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

THROUGHPUTER, INC., Plaintiff, v. AMAZON WEB SERVICES, INC., Defendant.	No. 1:22-cv-1095 **JURY TRIAL DEMANDED**

PLAINTIFF'S ORIGINAL COMPLAINT

Plaintiff ThroughPuter, Inc. ("ThroughPuter") files this Original Complaint and demand for jury trial seeking relief from patent infringement of the claims of U.S. Patent Nos. 11,347,556 and 11,385,934 (collectively, the "Asserted Patents") by Amazon Web Services, Inc. ("Defendant" or "AWS").

INTRODUCTION

1. A Field Programmable Gate Array (FPGA) is a specific type of microprocessor that can be reconfigured based on the tasks to be performed by it. In some situations, this reconfiguration takes place in between the performance of tasks by the FPGA. For example, the FPGAs can be configured and re-configured to provide hardware accelerators for data processing functions such as encoding/decoding, encryption/decryption or compression/decompression and then reconfigured to perform speech translation in addition to, *e.g.*, encryption and compression.

2. Defendant offers a service named Elastic Compute Cloud or EC2. On information and belief, EC2 was designed and developed by AWS and is the subject of patent applications filed by Amazon Technologies, Inc. ("ATI"), which is a holding company for intellectual property assets, including, for example patents. *See* Declaration of Scott Hayden in 6:20-cv-00272-ADA, Dkt. 23-4. EC2 is a web service designed to make web-scale cloud computing easier for developers. EC2

allows users to rent virtual computers to run their own computer applications providing them with flexibility to use the computing resources they need without incurring sunk costs in expensive hardware.

3. On November 30, 2016, AWS announced it was launching a new EC2 instance type called the F1. The EC2 F1 instances incorporate FPGA processors. Using these FPGA processors for the underlying dynamic parallel execution environment or architecture, AWS is able to achieve a power savings and increased processing speeds for it and the users of applications running on its cloud computing platform.

NATURE OF THE ACTION

4. This is an action for patent infringement arising under the Patent Laws of the United States, 35 U.S.C. § 1 *et seq*., including 35 U.S.C. § 271.

5. ThroughPuter brings this action to halt Amazon's infringement of its rights under the Patent Laws of the United States, 35 U.S.C. § 1 *et seq*., which arise under U.S. Patent No. 11,347,556 ("the '556 Patent"), which is attached hereto as Exhibit 1, and U.S. Patent No. 11,385,934 ("the '934 Patent"), which is attached hereto as Exhibit 2.

THE PARTIES

6. Plaintiff ThroughPuter, Inc. is a Delaware corporation having a principal place of business at 249 Richmond Road, Williamsburg, VA 23185. Plaintiff owns over 50 issued domestic and foreign patents and pending applications protecting its products, services and technologies. ThroughPuter's President, Mark Sandstrom, is the named inventor on each of such patents and applications.

7. Plaintiff is a member of The College of William & Mary's Launchpad business incubator, also known as the Miller Entrepreneurship Center.

8. Plaintiff has developed and continues to develop various products and services, including i) Estimator™, a machine learning Application Specific Processor (ASP)-as-a-service offering of the ThroughPuter Platform-as-a-Service (PaaS) project, and ii) Grafword™, an artificial intelligence ("AI") powered, graphical authentication service that is a pilot application of the Estimator™ machine learning microservice.

9. Estimator™ provides a streaming machine learning ("ML") microservice, to support AI applications in unpredictably changing operating environments. Estimator™ allows its prediction models and logic parameters to be adjusted continuously while the microservice is in operation, such that its predictions will stay tuned-in to the prevailing reality of its operating environment, as that may evolve over time or even change abruptly. An International Search Report recently conducted by the International Search Authority under the Patent Cooperation Treaty concluded that this technology is patentable. A beta version of the Estimator™ application programming interface is currently commercially available for 3rd party developer subscription at www.estimatorlab.com.

10. Grafword™ provides graphic based high-security password generation and authentication, such that the level of authentication challenge is adjusted according to a level of deviation of a given user's online session attributes from what is expected for the given username. Grafword™ thus provides both high security as well as, for the authentic users, convenience in online authentication. An International Search Report recently conducted by the International Search Authority under the Patent Cooperation Treaty also concluded that this technology is patentable. A beta version of Grafword™ is used for Estimator™ account creation and login: https://estimatorlab.com/landing.

11. Both Estimator™ and Grafword™ have the potential to change the space in which they are offered due to the advantages provided by ThroughPuter's claimed inventions such as increased throughput and latency.

12. On information and belief, Defendant is a Delaware corporation with a principal address of 410 Terry Avenue North, Seattle, Washington 98109 and is a subsidiary of Amazon.com, Inc. On information and belief, AWS has a regular and established place of business located at 11501 Alterra Pkwy, Austin, Texas 78758.

13. On information and belief, Defendant sold and offered to sell products and services throughout Texas, including in this judicial district, and introduced products and services that perform infringing methods or processes into the stream of commerce knowing that they would be sold in Texas and this judicial district.

14. Defendant is registered to do business in Texas and may be served via its registered agent at Corporation Services Company, located at 211 E. 7th Street, Suite 620, Austin, Texas 78701, and at its regular place of business, or anywhere that it may be found.

JURISDICTION AND VENUE

15. This is an action for patent infringement which arises under the Patent Laws of the United States, 35 U.S.C. § 1 *et seq*.

16. This Court has subject matter jurisdiction at least under 28 U.S.C. §§ 1331 and 1338.

17. Venue is proper in this District under 28 U.S.C. §§ 1391(b) and 1400(b) because Defendant has committed acts of infringement and has a regular and established place of business in this District.

18. This Court has personal jurisdiction over Defendant because: (i) Defendant is present within or has minimum contacts within the State of Texas and this judicial district;

(ii) Defendant has purposefully availed itself of the privileges of conducting business in the State of Texas and in this judicial district; and (iii) Plaintiff's cause of action arises directly from Defendant's business contacts and other activities in the State of Texas and in this judicial district, including through selling infringing products and services in this District and because Plaintiff's claims arise out of and relate to Defendant's acts of direct and indirect infringement in this District.

19. The Court's exercise of jurisdiction over Defendant would not offend traditional notions of fair play and substantial justice because Defendant has established at least the required minimum contacts with the forum.

20. Defendant maintains regular and established places of business throughout Texas and in this District, including offices housing, on information and belief, hundreds of employees at The Domain in Austin, a fulfillment center in the Metric Center in Austin, and numerous hubs throughout Texas and this District.

21. Defendant currently has 400+ positions available in Austin, Texas. *See* Exhibit 6 (search performed at https://www.amazon.jobs/ on October 11, 2022 with "Austin, TX" as the location and filtered by City: Austin and Team: Amazon Web Services).

22. For example, Defendant is currently looking for a "Software Development Engineer – FPGA Acceleration" in Austin, Texas. *See* Exhibit 7 (https://www.amazon.jobs/en/jobs/1917194/software-development-engineer-fpga-acceleration, last visited on October 11, 2022). That job summary sates, in part, "We are seeking an experienced FPGA Development Engineer to build the next generation of our cloud server platforms . . . Your results will optimize the hardware and software in our data centers including technologies such as AWS F1 which is an AWS EC2 instance used to deploy hardware accelerators in the cloud." *Id.*

As discussed *infra*, the AWS EC2 running F1 server instances is the basis for Plaintiff's allegations of infringement.

23. Defendant builds its own hardware at labs including one in Austin, Texas. *See* Exhibit 8 (https://www.aboutamazon.com/news/aws/take-a-look-inside-the-lab-where-aws-makes-custom-chips, last visited on October 11, 2022). Specifically, a specialist microelectronics company Defendant acquired in 2015, namely Annapurna Labs, designs and builds Defendant's own family of custom chips and accelerators. *See id* at 2. Employees are located in multiple locations around the world, including Tel Aviv, Israel; Toronto, Canada; and Austin, Texas. *See id.*

24. According to AWS, "AWS has its own family of custom chips and accelerators, with each new generation building on—and improving on—what came before. All the chips are designed and built by the Annapurna Labs team." *Id.*

25. Further, according to AWS, "[t]he [Annapurna Labs] team tests all new equipment in the mini data center before rolling it out for real." *Id.*

26. On information and belief, AWS commits acts of infringement in this District, including by, for example, making and using a mini data center that practices the inventions claimed in the '556 Patent and the '934 Patent, in which mini data center all new equipment is tested.

27. Defendant has substantial business contacts within this District and has purposefully availed itself of the privileges and benefits of the laws of the State of Texas.

BACKGROUND

28. This case involves ThroughPuter's patented cloud computing, computing acceleration and related technologies, which were developed starting in 2010.

29. As of that time, advancements in computing technologies had generally fallen into two categories. First, in the field conventionally referred to as high performance computing, the main objective has been maximizing the processing speed of a given computationally intensive program running on dedicated hardware. In this field, speed was traditionally achieved by assigning a combination of separate parallel processors to all work on the same program simultaneously. Second, in the field conventionally referred to as utility or cloud computing, the main objective has been to most efficiently share a given pool of computing hardware resources among a large number of client application programs.

30. Thus, in effect, one branch of computing innovation has been seeking to effectively use a large number of parallel processors to accelerate execution of a single application program by parallelizing its processing across a maximum possible number of processors. At the same time, another branch of computing innovation has been seeking to share a single pool of computing capacity among a large number of application programs to optimize utilization of processing capacity. The former efforts pursue maximizing processing speed of a single program. The latter efforts pursue maximizing utilization of processing capacity.

31. As of the time of ThroughPuter's pioneering patent filings starting in 2011, there had not been major synergies between the effort to increase processing speed of a single program on the one hand, and maximizing processing capacity utilization on the other. Indeed, pursuing one of these traditional objectives often happened at the expense of the other, placing the two objectives in tension with each other.

32. For instance, while dedicating an entire parallel processor based (super) computer to each individual application would increase processing speed of the individual programs, it would also cause severely sub-optimal computing resource utilization, as much of the capacity would be

idle much of the time. On the other hand, while seeking to improve utilization of computing systems by sharing their processing capacity among a number of applications would lead to enhanced resource utilization, it also tended to slow down processing of individual programs. As such, the overall cost-efficiency of computing was not improving as much as improvements toward either of the two traditional objectives would imply: traditionally, increases in processing speed came at the expense of system utilization efficiency, while overall system utilization efficiency maximization came at the expense of individual application processing speed.

33. The foregoing tension was exacerbated by the fact that even mainstream application performance requirements were increasingly exceeding the processing throughput achievable from a single CPU core, *e.g.*, due to the practical limits being reached on the CPU clock rates. This created an emerging requirement for intra-application parallel processing (at ever finer grades) even for mainstream programs in order to pursue satisfactory processing speeds, while these programs were to be increasingly hosted on cloud computing platforms where the processing resources would be shared among programs of multiple clients.

34. These internally parallelized and/or pipelined (concurrent) enterprise and web applications would ultimately be largely deployed on dynamically shared cloud computing infrastructure by entities such as Defendant using the technologies patented, pioneered and promoted by ThroughPuter.

35. Given the foregoing, there existed a need as of 2011 for supporting a large number of concurrent applications on dynamically shared parallel processing resource pools. This then-existing need for a new parallel computing architecture could be met by a system that enabled increasing the speed of executing application programs (including through execution of a given

application in parallel across multiple processor cores and/or using hardware accelerators) while at the same time improving the utilization of the available computing resources.

36. To address these problems, ThroughPuter developed hardware implemented dynamic resource management functionality including a scheduler, placer, inter-task communications and input/output system for use with multicore processor arrays dynamically shared among multiple concurrent applications, preferably to be deployed on FPGA processors. To that end, in this technology approach, the manycore processor array involves a fabric of reconfigurable cores that can be on-demand programmed to supply the needed mix or match of hardware accelerators. ThroughPuter's technology provided a cloud computing solution that enables accelerated processing speeds across multiple application programs while at the same time optimizing processing resource utilization.

37. ThroughPuter's hardware-based manycore fabric enables processing to be dynamically parallelized and hardware-accelerated, which results in optimized on-time processing throughput across the programs sharing an array of manycore processors. The effect for the client or end user of an accelerated service is increased processing speed and reduced cost base for delivering the application service, such that it becomes economically feasible for cloud service providers to support a range of performance intensive applications even without charge to end-users.

38. In recognition of ThroughPuter's innovative achievements, ThroughPuter's Mark Sandstrom was invited to speak at various high performance and cloud computing conferences starting in 2012. GigaOm selected ThroughPuter as one of eleven finalists to present at Launchpad 2012 in San Francisco, California. That same year, ThroughPuter was invited to present its

Dynamic Parallel Execution Environment (DPEE) based PaaS approach at the high performance computing start-up showcase at the Supercomputing 2012 conference ("SC12") in Provo, Utah.

39. ThroughPuter's novel manycore fabric led to industry recognition of ThroughPuter and its technology. For example, in January 2013, ThroughPuter was invited to publish an article in the Cloud Computing Journal, discussing PaaS based on novel manycore fabric. *See* Exhibit 9.

40. In addition, in September 2014, Mr. Sandstrom presented at the FPGAworld Conference in Stockholm, Sweden, on the topic of *Hardware Implemented Scheduler, Placer, Inter-Task Communications and IO System Functions for Manycore Processors Dynamically Shared among Multiple Applications*. *See* Exhibit 10.

41. By the time of the 2014 FPGAworld Conference, ThroughPuter had already been granted at least a dozen U.S. and United Kingdom patents protecting techniques enabling the advantages of its Dynamic Parallel Execution Environment™ (DPEE).

42. The following year, ThroughPuter was invited to present at the 2015 HPC Advisory Council Conference in Spain on the topic of executing multiple dynamically parallelized programs on dynamically shared cloud processors. A copy of the presentation is attached hereto as Exhibit 11.

AMAZON'S INFRINGING CLOUD COMPUTING ARCHITECTURE

43. Amazon's infringing cloud computing platform, discussed above at 1-42, is known as Elastic Compute Cloud ("EC2"). EC2 running F1 server instances (the "EC2/F1 Platform") is the technology service platform accused of infringement herein.

44. On information and belief, Defendant operates the EC2 cloud computing platform under the names of "Amazon Web Services" and "AWS."

45. ATI owns many registered trademarks for AWS and Amazon Web Services, including United States Trademark Registration Nos. 3576161, 5705338, and 6420129. *See* Exhibit 12; Exhibit 13; Exhibit 14.

46. On September 28, 2016, ATI filed U.S. Patent Application Serial No. 15/279,232 ("the '232 Application"). The corresponding patent, U.S. Patent No. 10,223,317 ("the '317 Patent"), is submitted herewith as Exhibit 3.

47. On September 29, 2016, ATI filed U.S. Patent Application Serial No. 15/280,624 ("the '624 Application"). The corresponding patent, U.S. Patent No. 10,282,330 ("the '330 Patent"), is submitted herewith as Exhibit 4.

48. On information and belief, Defendant's EC2/F1 Platform includes or has included the functionality described in the'317 and '330 Patents.

49. On information and belief, Defendant currently uses the technology described in its '317 and '330 Patents to manage at least a portion of its EC2/F1 Platform workloads.

50. The substantial identity between ATI's '317 Patent and ThroughPuter's invention can be appreciated from a side-by-side comparison of the patent claims granted to ATI and ThroughPuter. The column on the left shows claim 1 of the '934 Patent, which claims priority to applications dating back to 2013 and 2014, namely Provisional Application Serial No. 61/869,646, filed on August 23, 2013, Provisional Application Serial No. 61/934,747, filed on February 1, 2014, and Non-Provisional Application Serial No. 14/318,512, filed on June 27, 2014. The column on the right shows the text of independent claim 1 of ATI's '317 Patent, which application was filed on and claims a priority date of September 28, 2016. As can be appreciated from this side-by-side comparison, ATI obtained a patent on substantially the same technology taught by ThroughPuter's patent applications. However, ThroughPuter's '934 Patent is entitled to a priority date that is at

least two years earlier than ATI's earliest possible priority date, and possibly up to three years earlier.

ThroughPuter's U.S. Patent No. 11,385,934, Independent Claim 1	Amazon Technologies, Inc.'s U.S. Patent No. 10,223,317, Independent Claim 1
A configurable logic platform, the configurable logic platform comprising:	A configurable logic platform, the configurable logic platform comprising:
a physical interconnect for connecting the configurable logic platform to a processor;	a physical interconnect for connecting the configurable logic platform to a processor;
a reconfigurable logic region of an FPGA comprising logic blocks that are configured based on configuration data;	a reconfigurable logic region comprising logic blocks that are configured based on configuration data;
a configuration port for applying the configuration data to the reconfigurable logic region so that the reconfigurable logic region is configured based on configuration data;	a configuration port for applying the configuration data to the reconfigurable logic region so that the reconfigurable logic region is configured based on configuration data;
a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing only restricted access to the configuration port from the physical interconnect; and	a control plane function accessible via transactions of the physical interconnect, the control plane function in communication with the configuration port, the control plane function providing only restricted access to the configuration port from the physical interconnect; and
an interface function accessible via transactions of the physical interconnect, the interface function providing an interface to the reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the reconfigurable logic region from directly accessing the physical interconnect,	a data plane function accessible via transactions of the physical interconnect, the data plane function providing an interface to the reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the reconfigurable logic region from directly accessing the physical interconnect,
wherein the reconfiguration logic function is implemented in the reconfigurable logic region.	wherein the control plane function is implemented in the reconfigurable logic region.

51. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '317 Patent is patent eligible under 35 U.S.C. § 101.

52. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '317 Patent is novel under 35 U.S.C. § 102 in view of the references disclosed to and considered by the United States Patent & Trademark Office ("USPTO") during examination of the '232 Application.

53. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '317 Patent is non-obvious under 35 U.S.C. § 103 in view of the references Defendant disclosed to and subsequently were considered by the USPTO during examination of the '232 application.

54. Similarly, the substantial identity between ATI's '330 Patent and ThroughPuter's invention can be appreciated from a side-by-side comparison of the patent claims granted to ATI and ThroughPuter. The column on the left shows claim 1 of the '556 Patent, which claims priority to applications dating back to 2013 and 2014, namely Provisional Application Serial No. 61/869,646, filed on August 23, 2013, Provisional Application Serial No. 61/934,747, filed on February 1, 2014, and Non-Provisional Application Serial No. 14/318,512, filed on June 27, 2014. The column on the right shows the text of independent claim 1 of ATI's '330 Patent, which application was filed on and claims a priority date of September 29, 2016. As can be appreciated from this side-by-side comparison, ATI obtained a patent on substantially the same technology taught by ThroughPuter's patent applications. However, ThroughPuter's'556 Patent is entitled to a priority date that is at least two years earlier than ATI's earliest possible priority date, and possibly up to three years earlier.

ThroughPuter's U.S. Patent No. 11,347,556, Independent Claim 1	Amazon Technologies, Inc.'s U.S. Patent No. 10,282,330, Independent Claim 1
1. A configurable logic platform comprising:	1. A configurable logic platform comprising:
a physical interconnect for connecting the configurable logic platform to a processor;	a physical interconnect for connecting the configurable logic platform to a processor;
a first reconfigurable logic region comprising logic blocks that are configured based on first configuration data corresponding to the first reconfigurable logic region;	a first reconfigurable logic region comprising logic blocks that are configured based on configuration data corresponding to the first reconfigurable logic region;
a second reconfigurable logic region comprising logic blocks that are configured based on second configuration data corresponding to the second reconfigurable logic region;	a second reconfigurable logic region comprising logic blocks that are configured based on configuration data corresponding to the second reconfigurable logic region;

ThroughPuter's U.S. Patent No. 11,347,556, Independent Claim 1	Amazon Technologies, Inc.'s U.S. Patent No. 10,282,330, Independent Claim 1
a configuration port for applying the first and second configuration data to the first and second reconfigurable logic regions so that the first reconfigurable logic region is configured based on the first configuration data corresponding to the first reconfigurable logic region and the second reconfigurable logic region is configured based on the second configuration data corresponding to the second reconfigurable logic region;	a configuration port for applying the configuration data to the first and second reconfigurable logic regions so that the first reconfigurable logic region is configured based on the configuration data corresponding to the first reconfigurable logic region and the second reconfigurable logic region is configured based on the configuration data corresponding to the second reconfigurable logic region;
a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing restricted access to the configuration port from the physical interconnect;	a control plane function accessible via transactions of the physical interconnect, the control plane function in communication with the configuration port, the control plane function providing restricted access to the configuration port from the physical interconnect;
a first interface function accessible via transactions of the physical interconnect, the first interface function providing an interface to the first reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the first reconfigurable logic region from directly accessing the physical interconnect;	a first data plane function accessible via transactions of the physical interconnect, the first data plane function providing an interface to the first reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the first reconfigurable logic region from directly accessing the physical interconnect;
a second interface function accessible via transactions of the physical interconnect, the second interface function providing an interface to the second reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the second reconfigurable logic region from directly accessing the physical interconnect; and	a second data plane function accessible via transactions of the physical interconnect, the second data plane function providing an interface to the second reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the second reconfigurable logic region from directly accessing the physical interconnect; and
logic configured to apportion bandwidth of the physical interconnect among at least the first interface function and the second interface function.	arbitration logic configured to apportion bandwidth of the physical interconnect among at least the first data plane function and the second data plane function.

55. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '330 Patent is patent eligible under 35 U.S.C. § 101.

56. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '330 Patent is novel under 35 U.S.C. § 102 in view of the references disclosed to and considered by the USPTO during examination of the '624 Application.

57. On information and belief, Defendant believed at the time of filing and still believes that the subject matter claimed in the '330 Patent is non-obvious under 35 U.S.C. § 103 in view of the references Defendant disclosed to and subsequently were considered by the USPTO during examination of the '624 Application.

DEFENDANT'S INFRINGEMENT OF PLAINTIFF'S PATENTS

58. Defendant has been and is presently infringing, and will continue to infringe, the Asserted Patents in this District and elsewhere in the United States by, among other things, making, using, selling, offering for sale and/or importing the EC2/F1 Platform.

59. Defendant directly infringes the Asserted Patents pursuant to 35 U.S.C. § 271(a), either literally or under the doctrine of equivalents, or both.

60. Defendant also indirectly infringes the Asserted Patents by encouraging, instructing, directing, and requiring others, including its customers, purchasers, users, and developers, to use the configurable logic platform of the Asserted Patents, either literally or under the doctrine of equivalents, or both.

COUNT 1

(Direct Infringement of U.S. Patent No. 11,347,556)

61. ThroughPuter repeats and realleges each and every allegation contained above as though fully set forth herein.

62. On May 31, 2022, the USPTO duly and legally issued the '556 Patent, entitled "Configurable Logic Platform With Reconfigurable Processing Circuitry." A copy of the '556 Patent is attached as Exhibit 1.

63. Mark Sandstrom is the sole and true inventor of the ’556 Patent.

64. ThroughPuter owns all right, title, and interest to and in the ’556 Patent.

65. Claim 1 of the ’556 Patent is representative of the claims infringed by Defendant and recites:

> 1. A configurable logic platform comprising:
>
> a physical interconnect for connecting the configurable logic platform to a processor;
>
> a first reconfigurable logic region comprising logic blocks that are configured based on first configuration data corresponding to the first reconfigurable logic region;
>
> a second reconfigurable logic region comprising logic blocks that are configured based on second configuration data corresponding to the second reconfigurable logic region;
>
> a configuration port for applying the first and second configuration data to the first and second reconfigurable logic regions so that the first reconfigurable logic region is configured based on the first configuration data corresponding to the first reconfigurable logic region and the second reconfigurable logic region is configured based on the second configuration data corresponding to the second reconfigurable logic region;
>
> a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing restricted access to the configuration port from the physical interconnect;

a first interface function accessible via transactions of the physical interconnect, the first interface function providing an interface to the first reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the first reconfigurable logic region from directly accessing the physical interconnect;

a second interface function accessible via transactions of the physical interconnect, the second interface function providing an interface to the second reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the second reconfigurable logic region from directly accessing the physical interconnect; and

logic configured to apportion bandwidth of the physical interconnect among at least the first interface function and the second interface function.

66. On information and belief, Defendant has maintained, operated and administered and continues to maintain, operate, and administer data storage products and services that infringe one or more claims of the '556 Patent, including claim 1, literally or under the doctrine of equivalents.

67. Defendant put the inventions claimed by the '556 Patent into service (*i.e.*, used them); but for Defendant's actions, Defendant's products and services would never have been put into service. Defendant's acts complained herein caused those claimed-invention embodiments as a whole to perform, and Defendant's procurement of monetary and commercial benefit from it.

68. On information and belief, Defendant has directly infringed claim 1 of the '556 Patent by making, using, offering to sell, selling, and/or importing data storage products and services marketed as Amazon EC2, which can run F1 server instances.

69. On information and belief, the EC2/F1 Platform is a configurable logic platform. Defendant makes, uses, and sells cloud compute time on server "instances" of various types, including the EC2 F1 server instance with customer-configurable logic. "Amazon EC2 F1 instances use FPGAs [Field Programmable Gate Arrays] to enable delivery of custom hardware accelerations…Once your FPGA design is complete, you can register it as an Amazon FPGA Image (AFI), and deploy it to your F1 instance in just a few clicks. You can reuse your AFIs as many times as you like, and across as many F1 instances as you like."

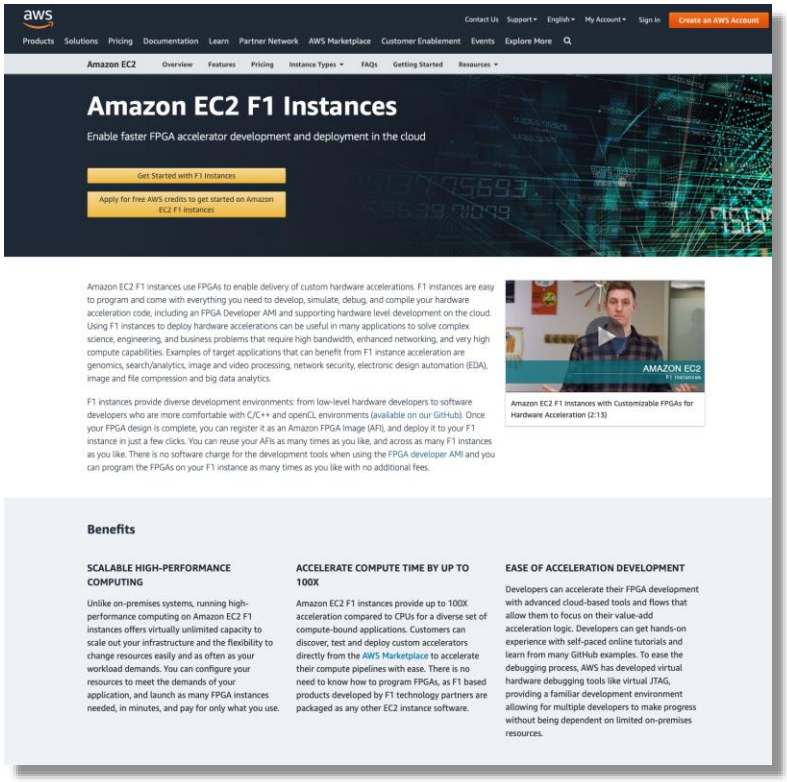

Exhibit 15 at 1. (https://aws.amazon.com/ec2/instance-types/f1/, last visited on October 11, 2022).

70. On information and belief, the EC2/F1 Platform comprises a physical interconnect for connecting the configurable logic platform to a processor. For example, a EC2/F1 Platform comprises a configurable logic platform containing a CPU (a processor) and one or more FPGAs.

A peripheral component interconnect express (PCIe) physical interconnect connects the FPGA(s) to the CPU:

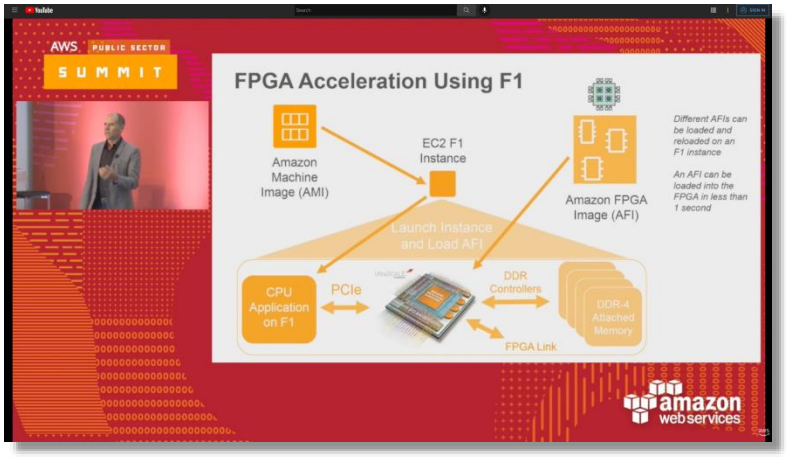

See https://youtu.be/IAQ-2C7yoQQ at 9:11 (last accessed on October 13, 2022).

71. On information and belief, the EC2/F1 Platform comprises a first reconfigurable logic region comprising logic blocks that are configured based on first configuration data corresponding to the first reconfigurable logic region; and a second reconfigurable logic region comprising logic blocks that are configured based on second configuration data corresponding to the second reconfigurable logic region. For example, each EC2 F1 server instance contains one or more FPGAs with customer-configurable logic regions, *i.e.*, "each with over 2 million customer-accessible FPGA programmable logic cells and over 5000 programmable DSP blocks."

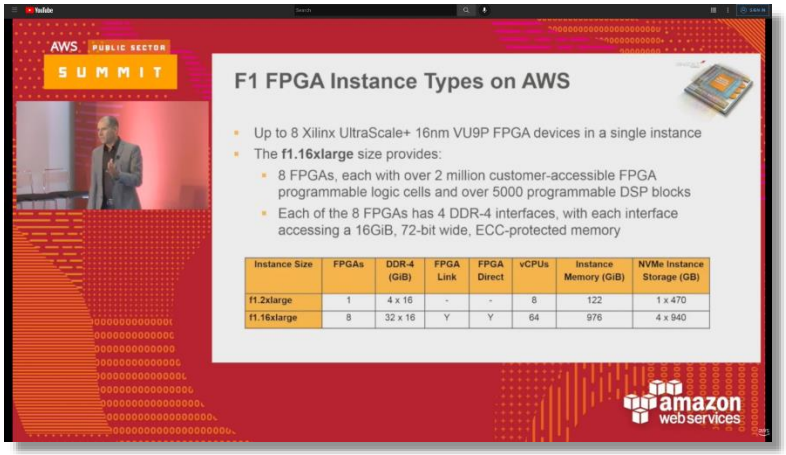

See https://youtu.be/IAQ-2C7yoQQ at 6:18 (last accessed on October 13, 2022).

72. On information and belief, the EC2/F1 Platform comprises a configuration port for applying the first and second configuration data to the first and second reconfigurable logic regions so that the first reconfigurable logic region is configured based on the first configuration data corresponding to the first reconfigurable logic region and the second reconfigurable logic region is configured based on the second configuration data corresponding to the second reconfigurable logic region. For example, AWS stores, in its cloud, configuration data for the configurable logic regions as Amazon FPGA Images (AFIs). "Different AFIs can be loaded and reloaded on an F1 instance. An AFI can be loaded into the FPGA in less than 1 second."

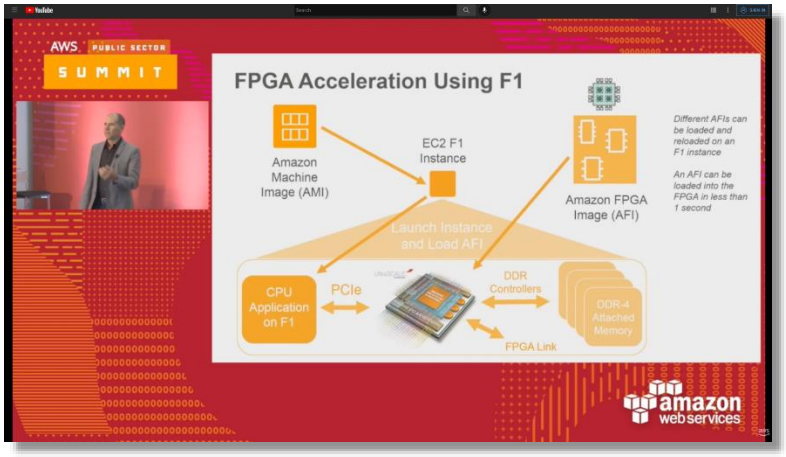

See https://youtu.be/IAQ-2C7yoQQ at 9:11 (last accessed on October 13, 2022).

73. Further, each EC2/F1 Platform wraps the reconfigurable logic regions in an "AWS FPGA Shell" that communicates with the reconfigurable logic regions via several ports, including an SDA AXI-L standardized interface port for applying AFI configuration data to the reconfigurable logic regions.

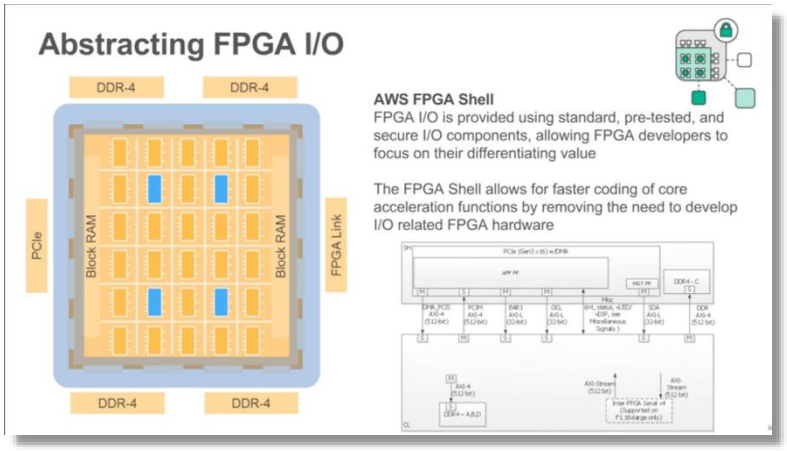

See https://youtu.be/R_Wxc8y7lb0 at 15:06 (last accessed on October 13, 2022).

74. On information and belief, the EC2/F1 Platform comprises a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing restricted access to the configuration port from the physical interconnect. For example, the MGT PF reconfiguration logic function is part of the PCIe interconnect logic and is thus accessible via transactions over the PCIe interconnect. The MGT PF reconfiguration logic function is also in communication with the SDA AXI-L master configuration port.

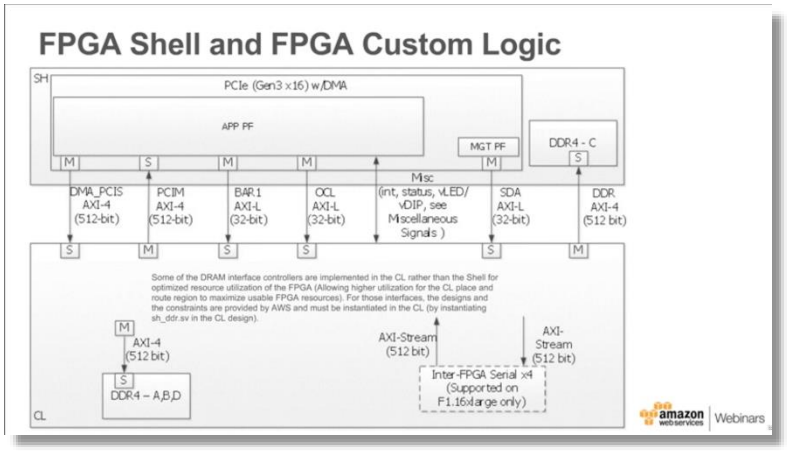

See https://youtu.be/R_Wxc8y7lb0 at 16:03 (last accessed on October 13, 2022).

75. Further, the management reconfiguration logic MGT PF (or MgmtPF) is a portion of the shell that restricts access to reconfiguration of the customer logic to AWS FPGA Mgmt Tools in Linux Userspace, which communicate through a protected Linux Kernel and over the PCIe physical interconnect to an FPGA.

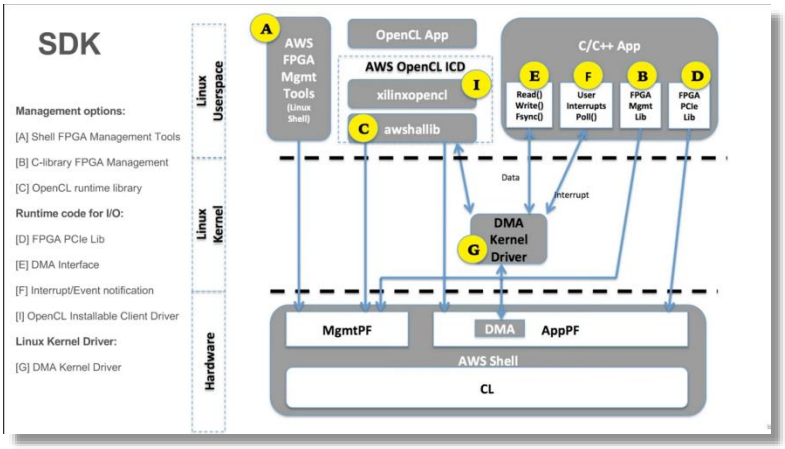

See https://youtu.be/R_Wxc8y7lb0 at 23:49 (last accessed on October 13, 2022).

76. Still further, the customer instantiates an AFI on an FPGA through the Linux Shell tools by specifying the image slot on the FPGA and AFI ID:

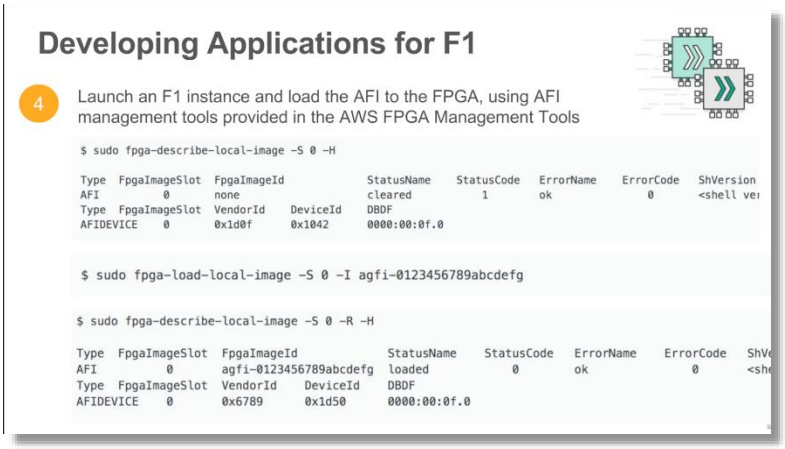

See https://youtu.be/R_Wxc8y7lb0 at 24:29 (last accessed on October 13, 2022).

77. On information and belief, the EC2/F1 Platform comprises a first interface function accessible via transactions of the physical interconnect, the first interface function

providing an interface to the first reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the first reconfigurable logic region from directly accessing the physical interconnect; and a second interface function accessible via transactions of the physical interconnect, the second interface function providing an interface to the second reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the second reconfigurable logic region from directly accessing the physical interconnect. For example, the customer logic regions of the FPGA are accessible via DMA (Direct Memory Access) transactions on the PCIe physical interconnect. "FPGA I/O is provided using standard, pre-tested, and secure I/O components, allowing FPGA developers to focus on their differentiating value." The customer logic regions must communicate through the AWS FPGA Shell and are prevented from accessing the PCIe physical interconnect directly.

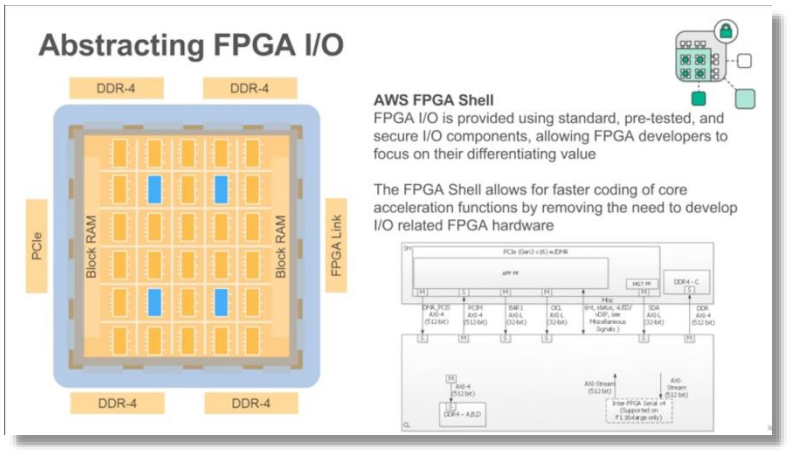

See https://youtu.be/R_Wxc8y7lb0 at 15:06 (last accessed on October 13, 2022).

78. On information and belief, the EC2/F1 Platform comprises logic configured to apportion bandwidth of the physical interconnect among at least the first interface function and the second interface function. For example, while the internal workings of the AWS FPGA Shell are not disclosed publicly to developers, ATI has, however, filed applications for patents that match closely to the description above, including granted patent US 11,182,320 B2, titled

"Configurable Logic Platform With Multiple Reconfigurable Regions" ("the '320 Patent"), which is attached hereto as Exhibit 5.

79. Figure 1 of the '320 Patent shows host logic 111 that includes a host interface 112 and data path functions 116A and 116B, communicating respectively with reconfigurable logic regions 140A and 140B:

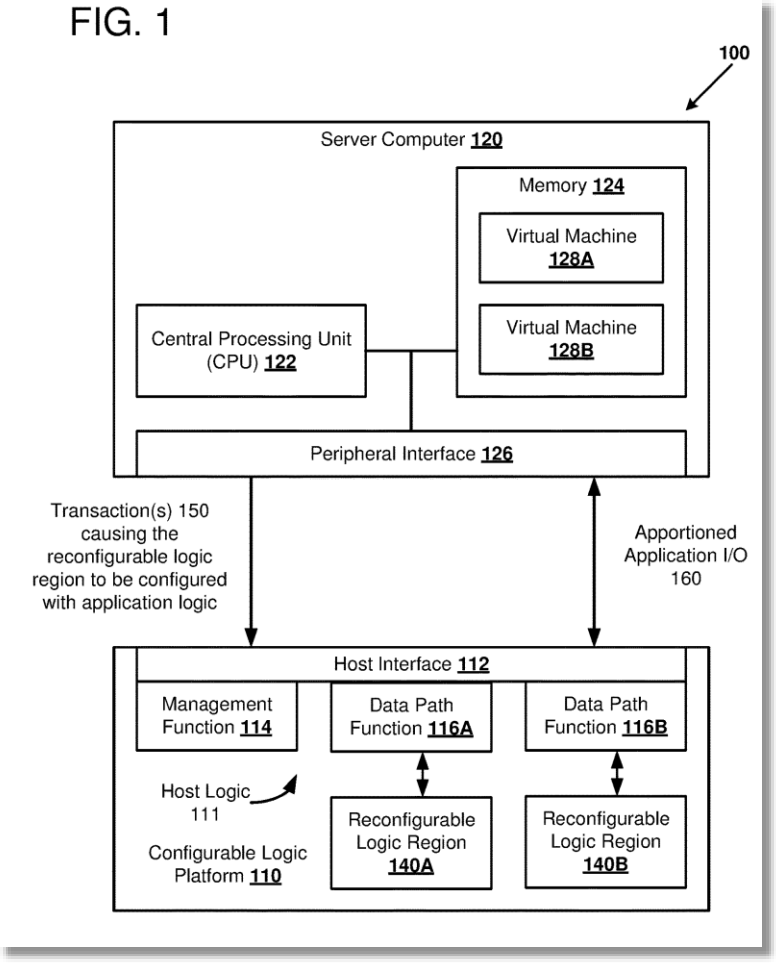

Exhibit 5, FIG. 1.

80. Figure 2 of the '320 Patent shows a more detailed block diagram, with a physical interconnect 230 that communicates with an interconnect interface 256 on a configurable hardware platform 210. Interconnect interface 256 contains arbitration logic 257, and is isolated

from the reconfigurable logic 240A and 240B by an address mapping layer 250, containing app functions 254A and 254B:

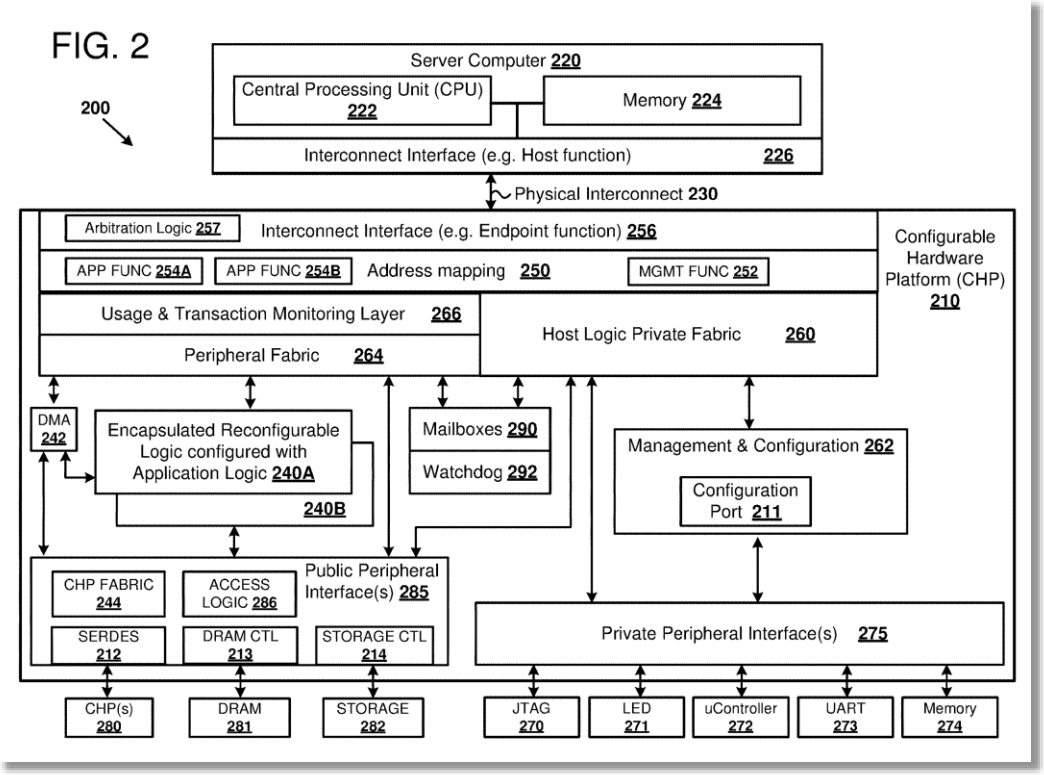

Exhibit 5, FIG. 2.

81. The description of Figures 1 and 2 in the ’320 Patent confirms that the interface and app/data path functions comprise logic that apportions bandwidth among the interface functions. For instance, “[t]he host interface 112 can enforce bandwidth, latency, size, and other quality of service factors for transactions over the physical interconnect. For example, the host interface 112 can apportion the outgoing bandwidth for transactions originating from the data path functions 116A-B and the management function 114.” ’320 Patent, col. 7, ll. 27-32.

82. Regarding Figure 2, “[t]he interconnect interface 256 can include arbitration logic 257 for apportioning bandwidth of the application functions 254A-B across the physical interconnect 230.” ’320 Patent, col. 13, ll. 57-59.

83. On information and belief, the functions described in the ’320 Patent are deployed in the AWS EC2/F1 Platform. For instance, Figure 1 above shows different virtual machines 128A and 128B deployed in the processor, for serving different customer application logic resident on the FPGA(s). The EC2 F1 server instances also support multiple virtual machines for hosting multiple customer applications that would compete for bandwidth, from 8 vCPUs on the single-FPGA server instance up to 64 vCPUs on the eight-FPGA server instance:

Product Details

Name	FPGAs	vCPUs	Instance Memory (GiB)	SSD Storage (GB)	Enhanced Networking	EBS Optimized	On-Demand Price/hr*	1-yr Reserved Instance Effective Hourly*	3-yr Reserved Instance Effective Hourly*
f1.2xlarge	1	8	122	470	Yes	Yes	$1.65	$1.06	$0.76
f1.4xlarge	2	16	244	940	Yes	Yes	$3.30	$2.12	$1.52
f1.16xlarge	8	64	976	4 x 940	Yes	Yes	$13.20	$8.50	$6.10

See Exhibit 15 at 6. (https://aws.amazon.com/ec2/instance-types/f1/, last visited on October 11, 2022.

84. The above allegations of infringement are preliminary and are therefore subject to change.

85. Defendant has caused Plaintiff damage by direct infringement of the claims of the ’556 Patent.

86. In accordance with 35 U.S.C. § 287, Defendant has had actual notice and knowledge of the ’556 Patent no later than the filing of this Complaint.

87. On information and belief, Defendant continues, without license, to make, use, import, market, offer for sale, and/or sell in the United States services or products that infringe the ’556 Patent.

88. Defendant has directly infringed and continues to directly infringe the ’556 Patent by engaging in acts constituting infringement under 35 U.S.C. § 271(a), including but not necessarily limited to one or more of making, using, selling and offering to sell, in this District

and elsewhere in the United States, and importing into the United States, the Elastic Compute Cloud (EC2) F1 Platform or components and services thereof.

89. Defendant's infringement of the '556 Patent has injured ThroughPuter in its business and property rights. ThroughPuter is entitled to recover monetary damages for the injuries arising from Defendant's infringement in an amount to be determined at trial.

90. Defendant's infringement of the '556 Patent has caused irreparable harm to ThroughPuter and will continue to cause such harm unless and until Defendant's infringing activities are enjoined by this Court.

91. Defendant has also had knowledge of the '556 Patent, and the way that its products infringe that patent, since at least the filing date of this Complaint.

92. Defendant's infringement of the '556 Patent is willful. Defendant continues to commit acts of infringement despite a high likelihood that its actions constitute infringement, and Defendant knew or should have known that its actions constituted an unjustifiably high risk of infringement. Defendant's continuing infringement after the filing of this Complaint is particularly egregious.

93. Defendant's infringement of the '556 Patent is exceptional and entitles ThroughPuter to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

94. The full extent of Defendant's infringement is not presently known to ThroughPuter. ThroughPuter makes this preliminary identification of infringing products and infringing claims in Count 1 without the benefit of discovery or claim construction in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

COUNT 2

(Indirect Infringement of U.S. Patent No. 11,347,556)

95. ThroughPuter repeats and realleges each and every allegation contained above as though fully set forth herein.

96. Defendant knew it was infringing the '556 Patent no later than the date it received this Complaint.

97. In addition to directly infringing the '556 Patent, as discussed above with respect to Count 1, Defendant knew or was willfully blind to the fact that it was inducing infringement of the '556 Patent under 35 U.S.C. § 271(b) by instructing, encouraging, directing, and requiring third parties, including its customers, to directly infringe by using the Accused Products in the United States.

98. Defendant knowingly and actively aided and abetted, encouraged, and contributed to the indirect infringement of the '556 Patent by instructing and encouraging its customers, purchasers, users, developers, vendors, partners, and manufacturers to meet the elements of the '556 Patent with the EC2/F1 Platform, as described above. Such instructions and encouragement included, but is not limited to, advising third parties to use the Defendant in an infringing manner through direct communications, training and support materials, and customer support regarding how to configure and use the EC2/F1 Platform, by advertising and promoting the use of the EC2/F1 Platform in an infringing manner, and distributing development kits, development Amazon Machine Images, tutorials, presentations, webinars, guidelines, videos, manuals, white papers, and trainings to third parties on how the EC2/F1 Platform must be used. *See, e.g.*, Exhibit 15 (EC2 F1 Instances website) (https://aws.amazon.com/ec2/instance-types/f1/); EC2 F1 Tutorial (https://youtu.be/ZYeNW9ZOqYw); EC2 F1 Tutorial (https://youtu.be/UQFAZyUgKls); EC2 F1

Tutorial (https://youtu.be/LvunFyJQGp4); EC2 F1 Tutorial (https://youtu.be/SwM4p3bbtFo); EC2 F1 Webinar (https://youtu.be/5sDVprory3c); AWS FPGA Development Kit (including all documentation on F1, internal FPGA interfaces, and compiler scripts for generating Amazon FPGA Images (AFIs) (https://github.com/aws/aws-fpga); AWS FPGA Developer AMI (Amazon Machine Image) (https://aws.amazon.com/marketplace/pp/prodview-gimv3gqbpe57k); Exhibit 16 (Amazon re: Invent 2019) (https://d1.awsstatic.com/events/reinvent/2019/Accelerate_applications_using_Amazon_EC2_F1_FPGA_instances_CMP314.pdf); Exhibit 17 (Webinar: Deep Dive on Amazon EC2 F1 Instance) (https://pages.awscloud.com/GLB-WBNR-AWS-OTT-2017-05-16-Deep-Dive-on-Amazon-EC2-F1-Instance_RegPage.html); Exhibit 18 (Developing Cloud Scale FPGA Accelerations Using AWS F1) (https://h2rc.cse.sc.edu/2017/slides/amazon.pdf).

99. The above allegations of infringement are preliminary and are therefore subject to change.

100. Defendant's indirect infringement of the '556 Patent has injured and continues to injure Plaintiff in an amount to be proven at trial, but not less than a reasonable royalty. Defendant's indirect infringement of the '556 Patent has caused and is continuing to cause damage and irreparable injury to Plaintiff, and Plaintiff will continue to suffer damage and irreparable injury unless and until that infringement is enjoined by this Court.

101. Defendant has also had knowledge of the '556 Patent, and the way that its products infringe that patent, since at least the filing date of this Complaint.

102. Defendant's infringement of the '556 Patent is willful. Defendant continues to commit acts of infringement despite a high likelihood that its actions constitute infringement, and Defendant knew or should have known that its actions constituted an unjustifiably high risk of

infringement. Defendant's continuing infringement after the filing of this Complaint is particularly egregious.

103. Defendant's infringement of the '556 Patent is exceptional and entitles ThroughPuter to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

104. The full extent of Defendant's infringement is not presently known to ThroughPuter. ThroughPuter makes this preliminary identification of infringing products and infringing claims in Count 2 without the benefit of discovery or claim construction in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

COUNT 3

(Direct Infringement of U.S. Patent No. 11,385,934)

105. ThroughPuter repeats and realleges each and every allegation contained above as though fully set forth herein.

106. On July 12, 2022, the USPTO duly and legally issued the '934 Patent, entitled "Configurable Logic Platform With Reconfigurable Processing Circuitry." A copy of the '934 Patent is attached as Exhibit 2.

107. Mark Sandstrom is the sole and true inventor of the '934 Patent.

108. ThroughPuter owns all right, title, and interest to and in the '934 Patent.

109. Claim 1 of the '934 Patent is representative of the claims infringed by Defendant and recites:

> A configurable logic platform, the configurable logic platform comprising:
>
> a physical interconnect for connecting the configurable logic platform to a processor;

a reconfigurable logic region of an FPGA comprising logic blocks that are configured based on configuration data;

a configuration port for applying the configuration data to the reconfigurable logic region so that the reconfigurable logic region is configured based on configuration data;

a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing only restricted access to the configuration port from the physical interconnect; and

an interface function accessible via transactions of the physical interconnect, the interface function providing an interface to the reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the reconfigurable logic region from directly accessing the physical interconnect, wherein the reconfiguration logic function is implemented in the reconfigurable logic region.

110. On information and belief, Defendant has maintained, operated and administered and continues to maintain, operate and administer data storage products and services that infringe one or more claims of the ’934 Patent, including claim 1, literally or under the doctrine of equivalents.

111. Defendant put the inventions claimed by the ’934 Patent into service (*i.e.*, used them); but for Defendant’s actions, Defendant’s products and services would never have been put into service. Defendant’s acts complained herein caused those claimed-invention embodiments as

a whole to perform, and Defendant's procurement of monetary and commercial benefit from it.

112. On information and belief, Defendant has directly infringed claim 1 of the '934 Patent by making, using, offering to sell, selling, and/or importing the EC2/F1 Platform.

113. On information and belief, the EC2/F1 Platform is a configurable logic platform. Defendant makes, uses, and sells cloud compute time on server "instances" of various types, including the EC2 F1 server instance with customer-configurable logic. "Amazon EC2 F1 instances use FPGAs [Field Programmable Gate Arrays] to enable delivery of custom hardware accelerations…Once your FPGA design is complete, you can register it as an Amazon FPGA Image (AFI), and deploy it to your F1 instance in just a few clicks. You can reuse your AFIs as many times as you like, and across as many F1 instances as you like."

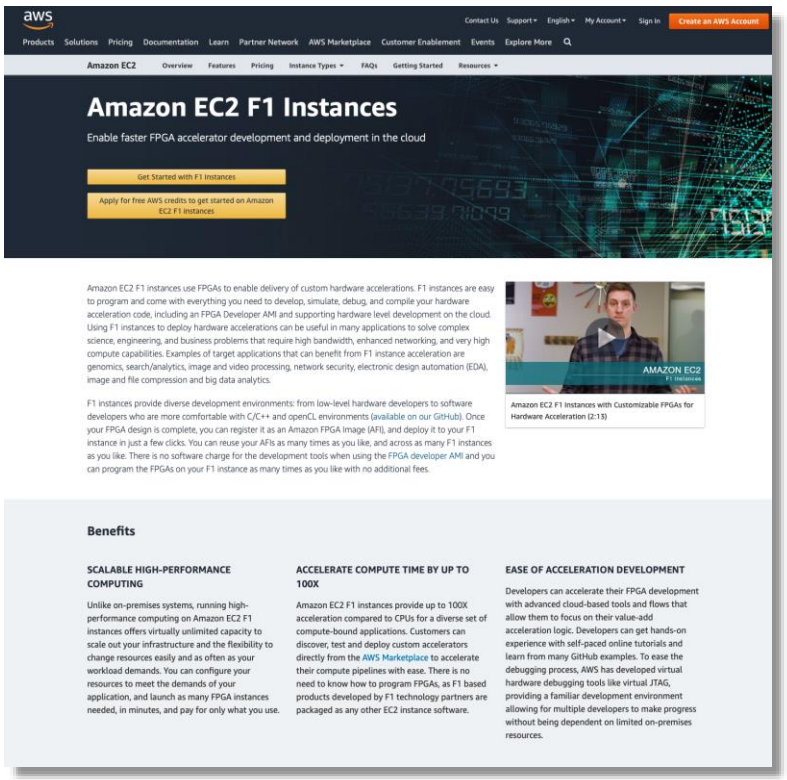

See Exhibit 15 at 1.

114. On information and belief, the EC2/F1 Platform comprises a physical interconnect for connecting the configurable logic platform to a processor. For example, a EC2/F1 Platform comprises a configurable logic platform containing a CPU (a processor) and one or more FPGAs. A PCIe physical interconnect connects the FPGA(s) to the CPU:

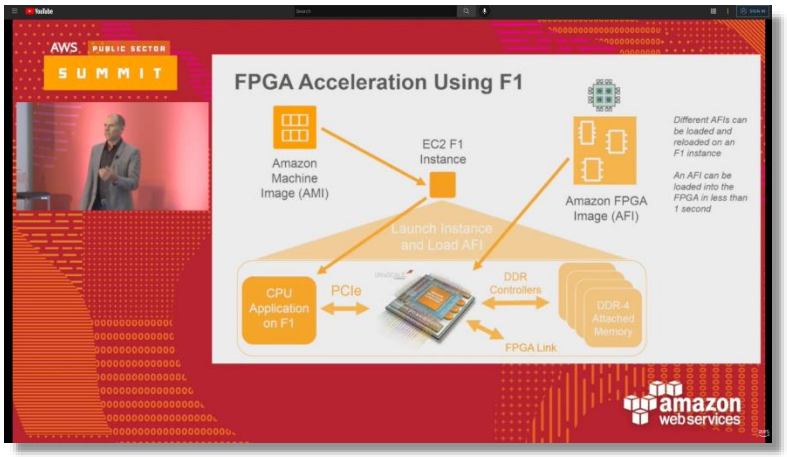

See https://youtu.be/IAQ-2C7yoQQ at 9:11 (last accessed on October 13, 2022).

115. On information and belief, the EC2/F1 Platform comprises a reconfigurable logic region of an FPGA comprising logic blocks that are configured based on configuration data. For example, each EC2 F1 server instance contains one or more FPGAs with customer-configurable logic regions, *i.e.*, "each with over 2 million customer-accessible FPGA programmable logic cells and over 5000 programmable DSP blocks."

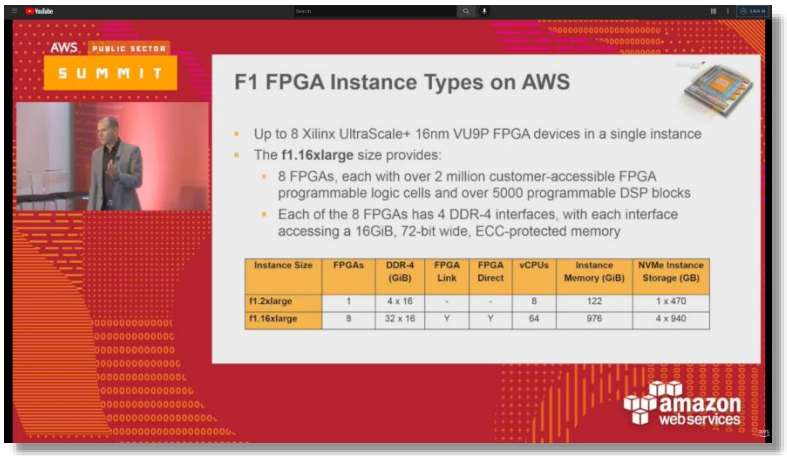

See https://youtu.be/IAQ-2C7yoQQ at 6:18 (last accessed on October 13, 2022).

116. On information and belief, the EC2/F1 Platform comprises a configuration port for applying the configuration data to the reconfigurable logic region so that the reconfigurable logic region is configured based on configuration data. For example, AWS stores, in its cloud, configuration data for the configurable logic regions as Amazon FPGA Images (AFIs). "Different AFIs can be loaded and reloaded on an F1 instance. An AFI can be loaded into the FPGA in less than 1 second."

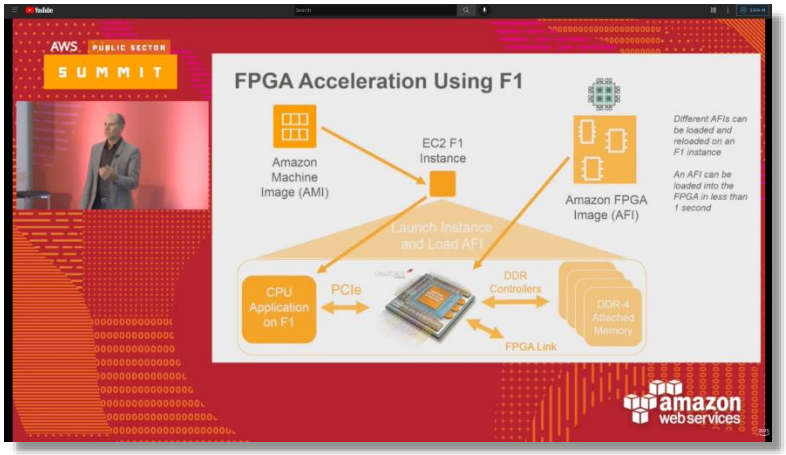

See https://youtu.be/IAQ-2C7yoQQ at 9:11 (last accessed on October 13, 2022).

117. Further, each EC2 F1 server instance wraps the reconfigurable logic regions in an "AWS FPGA Shell" that communicates with the reconfigurable logic regions via several ports, including an SDA AXI-L standardized interface port for applying AFI configuration data to the reconfigurable logic regions.

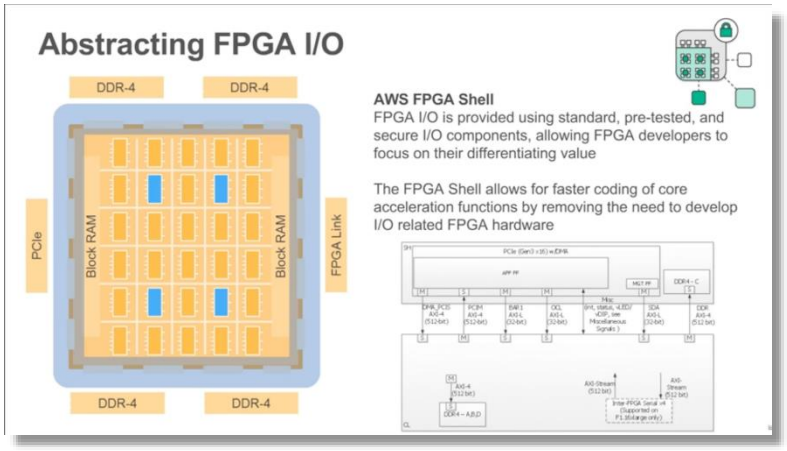

See https://youtu.be/R_Wxc8y7lb0 at 15:06 (last accessed on October 13, 2022.

118. On information and belief, the EC2/F1 Platform comprises a reconfiguration logic function accessible via transactions of the physical interconnect, the reconfiguration logic function in communication with the configuration port, the reconfiguration logic function providing only restricted access to the configuration port from the physical interconnect. For example, the MGT PF reconfiguration logic function is part of the PCIe interconnect logic and is thus accessible via transactions over the PCIe interconnect. The MGT PF reconfiguration logic function is also in communication with the SDA AXI-L master configuration port.

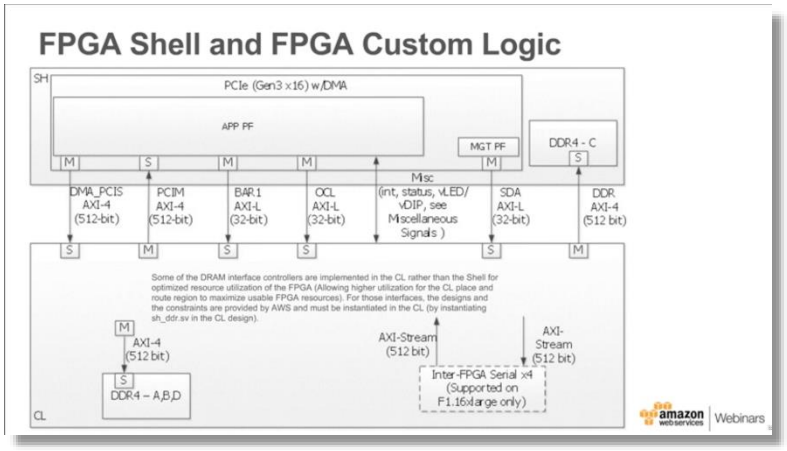

See https://youtu.be/R_Wxc8y7lb0 at 16:03 (last accessed on October 13, 2022).

119. Further, the management reconfiguration logic MGT PF (or MgmtPF) is a portion of the shell that restricts access to reconfiguration of the customer logic to AWS FPGA Mgmt Tools in Linux Userspace, which communicate through a protected Linux Kernel and over the PCIe physical interconnect to an FPGA.

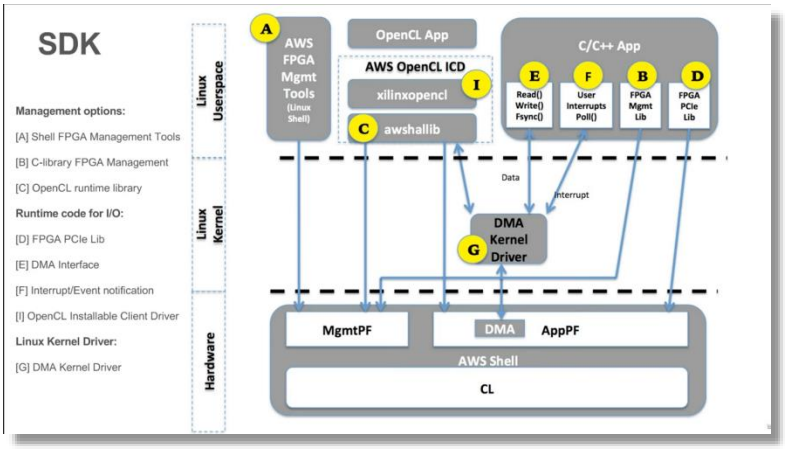

See https://youtu.be/R_Wxc8y7lb0 at 23:49 (last accessed on October 13, 2022).

120. Still further, the customer instantiates an AFI on an FPGA through the Linux Shell tools by specifying the image slot on the FPGA and AFI ID:

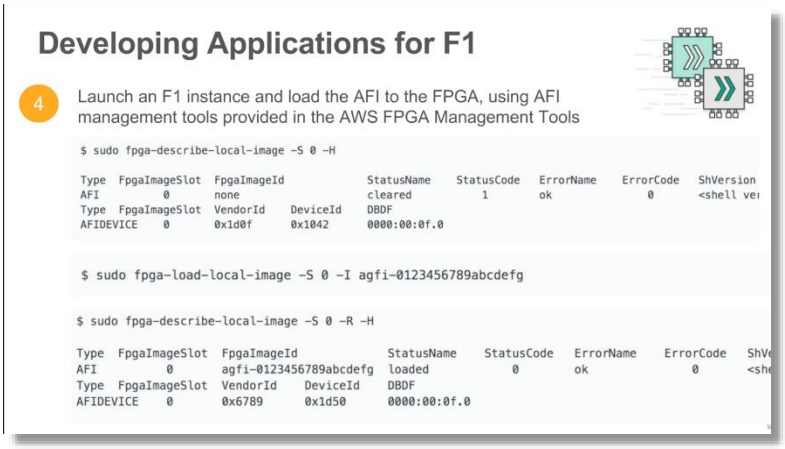

See https://youtu.be/R_Wxc8y7lb0 at 24:29 (last accessed on October 13, 2022).

121. On information and belief, the EC2/F1 Platform comprises an interface function accessible via transactions of the physical interconnect, the interface function providing an

interface to the reconfigurable logic region which allows information to be transmitted over the physical interconnect and prevents the reconfigurable logic region from directly accessing the physical interconnect, wherein the reconfiguration logic function is implemented in the reconfigurable logic region. For example, the customer logic regions of the FPGA are accessible via DMA (Direct Memory Access) transactions on the PCIe physical interconnect. "FPGA I/O is provided using standard, pre-tested, and secure I/O components, allowing FPGA developers to focus on their differentiating value." The customer logic regions must communicate through the AWS FPGA Shell and are prevented from accessing the PCIe physical interconnect directly.

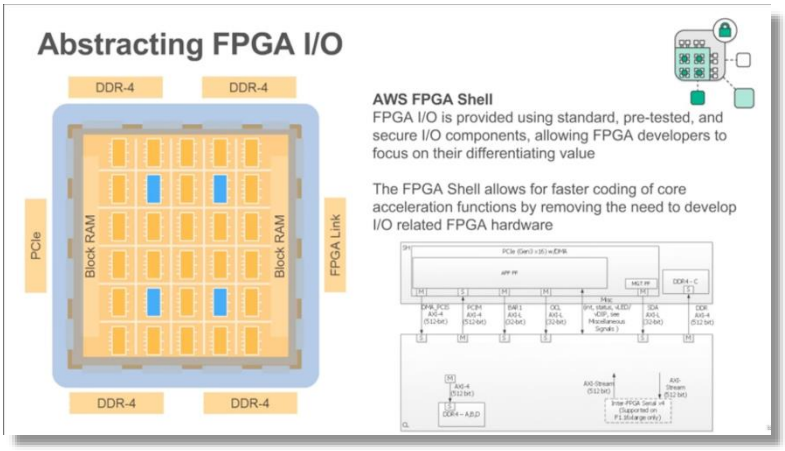

See https://youtu.be/R_Wxc8y7lb0 at 15:06 (last accessed on October 13, 2022).

122. Further, the FPGAs used in the EC2/F1 Platform are standard Xilinx parts with PCIe and DDR-4 physical interfaces, but must have sections of their reconfigurable logic regions programmed by AWS to include the AWS FPGA Shell. As stated by AWS, because they have programmed the shell it removes "the need to develop I/O related FPGA hardware," because Amazon has abstracted the interface logic and pre-programmed it. *See* https://youtu.be/R_Wxc8y7lb0 at 15:06 (last accessed on October 13, 2022).

123. These allegations of infringement are preliminary and are therefore subject to change.

124. Defendant has caused Plaintiff damage by direct infringement of the claims of the ’934 Patent.

125. In accordance with 35 U.S.C. § 287, Defendant has had actual notice and knowledge of the ’934 Patent no later than the filing of this Complaint.

126. On information and belief, Defendant continues, without license, to make, use, import, market, offer for sale, and/or sell in the United States services or products that infringe the ’934 Patent.

127. Defendant has directly infringed and continues to directly infringe the ’934 Patent by engaging in acts constituting infringement under 35 U.S.C. § 271(a), including but not necessarily limited to one or more of making, using, selling and offering to sell, in this District and elsewhere in the United States, and importing into the United States, the Elastic Compute Cloud (EC2) F1 Platform or components and services thereof.

128. Defendant’s infringement of the ’934 Patent has injured ThroughPuter in its business and property rights. ThroughPuter is entitled to recover monetary damages for the injuries arising from Defendant’s infringement in an amount to be determined at trial. Defendant’s infringement of the ’934 Patent has caused irreparable harm to ThroughPuter and will continue to cause such harm unless and until Defendant’s infringing activities are enjoined by this Court.

129. Defendant has also had knowledge of the ’934 Patent, and the way that its products infringe that patent, since at least the filing date of this Complaint.

130. Defendant’s infringement of the ’934 Patent is willful. Defendant continues to commit acts of infringement despite a high likelihood that its actions constitute infringement, and Defendant knew or should have known that its actions constituted an unjustifiably high risk of

infringement. Defendant's continuing infringement after the filing of this Complaint is particularly egregious.

131. Defendant's infringement of the '934 Patent is exceptional and entitles ThroughPuter to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

132. The full extent of Defendant's infringement is not presently known to ThroughPuter. ThroughPuter makes this preliminary identification of infringing products and infringing claims in Count 3 without the benefit of discovery or claim construction in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

COUNT 4

(Indirect Infringement of U.S. Patent No. 11,385,934)

133. ThroughPuter repeats and realleges each and every allegation contained above as though fully set forth herein.

134. Defendant knew it was infringing the '934 Patent no later than the date it received this Complaint.

135. In addition to directly infringing the '934 Patent, as discussed above with respect to Count 3, Defendant also knew or was willfully blind to the fact that it was inducing infringement of the '934 Patent under 35 U.S.C. § 271(b) by instructing, encouraging, directing, and requiring third parties, including its customers, to directly infringe by using the Accused Products in the United States.

136. Defendant knowingly and actively aided and abetted, encouraged, and contributed to the indirect infringement of the '934 Patent by instructing and encouraging its customers, purchasers, users, developers, vendors, partners, and manufacturers to meet the elements of the

'934 Patent with the EC2/F1 Platform, as described above. Such instructions and encouragement included, but is not limited to, advising third parties to use the Defendant in an infringing manner through direct communications, training and support materials, and customer support regarding how to configure and use the EC2/F1 Platform, by advertising and promoting the use of the EC2/F1 Platform in an infringing manner, and distributing development kits, development Amazon Machine Images, tutorials, presentations, webinars, guidelines, videos, manuals, white papers, and trainings to third parties on how the EC2/F1 Platform must be used. *See*, *e.g.*, Exhibit 15 (EC2 F1 Instances website) (https://aws.amazon.com/ec2/instance-types/f1/); EC2 F1 Tutorial (https://youtu.be/ZYeNW9ZOqYw); EC2 F1 Tutorial (https://youtu.be/UQFAZyUgKls); EC2 F1 Tutorial (https://youtu.be/LvunFyJQGp4); EC2 F1 Tutorial (https://youtu.be/SwM4p3bbtFo); EC2 F1 Webinar (https://youtu.be/5sDVprory3c); AWS FPGA Development Kit (including all documentation on F1, internal FPGA interfaces, and compiler scripts for generating Amazon FPGA Images (AFIs) (https://github.com/aws/aws-fpga); AWS FPGA Developer AMI (Amazon Machine Image) (https://aws.amazon.com/marketplace/pp/prodview-gimv3gqbpe57k); Exhibit 16 (Amazon re: Invent 2019) (https://d1.awsstatic.com/events/reinvent/2019/Accelerate_applications_using_Amazon_EC2_F1_FPGA_instances_CMP314.pdf); Exhibit 17 (Webinar: Deep Dive on Amazon EC2 F1 Instance) (https://pages.awscloud.com/GLB-WBNR-AWS-OTT-2017-05-16-Deep-Dive-on-Amazon-EC2-F1-Instance_RegPage.html); Exhibit 18 (Developing Cloud Scale FPGA Accelerations Using AWS F1) (https://h2rc.cse.sc.edu/2017/slides/amazon.pdf).

137. The above allegations of infringement are preliminary and are therefore subject to change.

138. Defendant's indirect infringement of the '934 Patent has injured and continues to injure Plaintiff in an amount to be proven at trial, but not less than a reasonable royalty. Defendant's indirect infringement of the '934 Patent has caused and is continuing to cause damage and irreparable injury to Plaintiff, and Plaintiff will continue to suffer damage and irreparable injury unless and until that infringement is enjoined by this Court.

139. Defendant has also had knowledge of the '934 Patent, and the way that its products infringe that patent, since at least the filing date of this Complaint.

140. Defendant's infringement of the '934 Patent is willful. Defendant continues to commit acts of infringement despite a high likelihood that its actions constitute infringement, and Defendant knew or should have known that its actions constituted an unjustifiably high risk of infringement. Defendant's continuing infringement after the filing of this Complaint is particularly egregious.

141. Defendant's infringement of the '934 Patent is exceptional and entitles ThroughPuter to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

142. The full extent of Defendant's infringement is not presently known to ThroughPuter. ThroughPuter makes this preliminary identification of infringing products and infringing claims in Count 4 without the benefit of discovery or claim construction in this action, and expressly reserves the right to augment, supplement, and revise its identifications based on additional information obtained through discovery or otherwise.

PRAYER FOR RELIEF

WHEREFORE, ThroughPuter respectfully requests that the Court enter judgment against Defendant as follows:

A. An adjudication that Defendant has directly and/or indirectly infringed one or more

claims of the ’934 Patent;

B. Entry of judgment declaring that Defendant’s infringement of the ’934 Patent is willful;

C. An adjudication that Defendant has directly and/or indirectly infringed one or more claims of the ’556 Patent;

D. Entry of judgment declaring that Defendant’s infringement of the ’556 Patent is willful;

E. An order permanently enjoining Defendant from further infringement of the ’934 Patent;

F. An order permanently enjoining Defendant from further infringement of the ’556 Patent;

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

H. An order that the damages award be increased up to three times the actual amount assessed, pursuant to 35 U.S.C. § 284;

I. An award to ThroughPuter of its costs, pre- and post-judgment interest, and reasonable expenses to the fullest extent permitted by law;

J. A declaration that this case is exceptional pursuant to 35 U.S.C. § 285, and an award of attorneys’ fees and costs; and

K. An award to ThroughPuter of such other and further relief as this Court deems just and proper.

DEMAND FOR JURY TRIAL

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

DATED: October 27, 2022

Respectfully submitted,

/s/ Melissa R. Smith
Melissa R. Smith
State Bar No. 24001351
GILLAM & SMITH, LLP
303 South Washington Avenue,
Marshall, Texas 75670
Telephone: (903) 934-8450
Facsimile: (903) 934-9257
melissa@gillamsmithlaw.com

Gardella Grace P.A.
W. Cook Alciati (to be admitted *pro hac vice*)
80 M Street SE, 1st Floor
Washington D.C. 20003
T: 703-721-8379
calciati@gardellagrace.com

Michael Dorfman (to be admitted *pro hac vice*)
2502 North Clark Street, Suite 222
Chicago, Illinois 60614
T: 773-755-4942
mdorfman@gardellagrace.com

Attorneys for Plaintiff ThroughPuter, Inc.