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

Concurrent Ventures, LLC and XtreamEdge, Inc.,	§	
	§	
Plaintiffs,	§	
v.	§	**CIVIL ACTION NO. 1:24-cv-335**
	§	
Advanced Micro Devices, Inc. and Pensando Systems, Inc.,	§	**JURY TRIAL DEMANDED**
	§	
Defendants.	§	

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiffs Concurrent Ventures, LLC ("Concurrent Ventures") and XtreamEdge, Inc. ("XtreamEdge") (collectively, "Plaintiffs") file this Complaint against Defendants Advanced Micro Devices, Inc. ("AMD") and Pensando Systems, Inc. ("Pensando") (collectively, "Defendants"), and in support thereof alleges as follows:

I.
NATURE OF THE ACTION

1. Data storage, manipulation, and use serves as the foundation for the global economy and have all increased exponentially over the past years. Data is integral to how our society operates—how we communicate, interact, and engage with entertainment on personal computing devices; how we do business; and how we power the artificial intelligence revolution.

2. As data becomes even more firmly entrenched in society, data processing and data functions require more resources. The traditional solution to this problem is to add more processors, in the form of additional Central Processing Units ("CPUs"). CPUs are general purpose hardware that implement functionality in software. This allows for a flexible approach, but these software solutions are slower and more resource-intensive than hardware-based solutions, and do not work well at scale.

3. Adding more processors is not sustainable. Additional CPUs lead to higher capital expenditures and operational expenditures, lower flexibility, and increased latency. Simply using CPUs worked as a stopgap for years because the power of CPUs increased at a sufficient pace in accordance with Moore's Law.[1] But as we near the end of Moore's Law, and as data use and consumption continues to increase, there is a need for more ***efficient*** data processing.

4. The inventors of the Asserted Patents addressed this need through a number of innovations that increase the efficiency of communications between CPUs and other network components. Plaintiffs' innovations, among other benefits, solves long-standing problems by offloading data-centric computations, which are estimated to be between 31-83% of the workload,[2] to more efficient units designed to handle data computations. By offloading the networking, communication, and other tasks from the CPU, the CPU is freed up to focus on what it does best.

5. Plaintiffs patented the innovations necessary to create and use this specialized unit, which it referred to as a Stream Processing Unit ("SPU"). Although it was the first, it was not the last. Pensando Systems Inc. pursued similar solutions as XtreamEdge—years after XtreamEdge—but called its product a Data Processing Unit ("DPU").[3] AMD acquired Pensando in 2022, leading to the integration of Pensando's DPU-enabled systems into AMD's networking, computing, and security services platforms.

---

[1] According to Moore's Law, the number of transistors in a dense integrated circuit doubles about every two years. https://en.wikipedia.org/wiki/Moore%27s_law; https://www.technologyreview.com/2020/02/24/905789/were-not-prepared-for-the-end-of-moores-law.

[2] https://www.sdxcentral.com/articles/interview/amd-shares-vision-for-pensando-smart-switch-dpu/2023/04.

[3] https://techcrunch.com/2022/04/04/amd-to-acquire-data-center-optimization-startup-pensando-for-2b.

6. Defendants and industry leaders have all noted the importance of DPUs in the evolution of data processing and its role as the third leg of processing, alongside CPUs and Graphics Processing Units ("GPUs"). *See, e.g.*, https://www.sdxcentral.com/articles/interview/amd-shares-vision-for-pensando-smart-switch-dpu/2023/04 ("'If you look at the journey that AMD is on in the data center, it's really to look at how to go about building the industry's highest performing adaptive computing portfolio,' Jiandani said. 'I think DPUs are the most efficient way to deal with things that need to have data processing at the edge.'"); https://www.techradar.com/pro/future-servers-could-have-a-shared-dpu-could-the-next-decade-see-a-rise-in-socket-heterogeneity (predicting that "there is a world where all hyperscalers will leverage servers with DPUs" and that "AMD is ideally positioned to be the one-stop shop for compute."). Robert Hormuth, AMD Vice President of Architecture & Strategy, stated that "[m]odern infrastructure is moving to offload much of the OS and hypervisor software stack to the DPU/SmartNIC. DPUs like the AMD Pensando™ enable DPU sharing across multiple servers. Single-socket servers are ideal for this new deployment model - connecting multiple single-socket servers to one DPU." https://www.amd.com/en/solutions/data-center/insights/myths-and-urban-legends-about-dual-socket-servers.html. AMD has also recognized the importance of adaptive computing engines—such as those implementing DPUs—in its artificial intelligence ("AI") space. https://www.nextplatform.com/2023/05/03/amd-says-ai-is-the-number-one-priority-right-now/?mc_cid=c589551dac&mc_eid=57a4c0e31b.

7. Defendants' refusal to acknowledge Plaintiffs' patents and offer fair compensation for their use violates the patent laws and undermines the effort that Plaintiffs took to develop their innovations. Plaintiffs feel they have no recourse but to file this action to stop Defendants' unauthorized use of Plaintiffs' patents. Plaintiffs invented something groundbreaking and

Defendants are taking advantage of those innovations without permission. Accordingly, Plaintiffs brings this action under the patent laws, 35 U.S.C. § 1 et seq., in order to stop Defendants' willful infringement of U.S. Patent Nos. 8,924,596; 10,873,753; 10,985,943; 10,944,634; 9,529,767 (collectively, the "Asserted Patents").

II.
PARTIES

8. XtreamEdge, Inc. is a corporation organized and existing under the laws of the State of Delaware, with its principal place of business in Johns Creek, Georgia. XtreamEdge is the assignee of the Asserted Patents.

9. Concurrent Ventures, LLC is a limited liability corporation organized and existing under the laws of the State of Georgia, with its principal place of business in Johns Creek, Georgia. Concurrent Ventures is the exclusive licensee to the Asserted Patents and has all substantial rights to the Asserted Patents.

10. Advanced Micro Devices, Inc. is a corporation organized and existing under the laws of Delaware that maintains an established place of business at 7171 Southwest Parkway, Austin, Texas 78735.

11. Pensando Systems, Inc. is a corporation organized and existing under the laws of Delaware that, on information and belief, maintains an established place of business at 7171 Southwest Parkway, Austin, Texas 78735.

III.
ASSERTED PATENTS

12. U.S. Patent No. 8,924,596 (the "'596 Patent"), titled "System and Method for Dividing and Synchronizing a Processing Task Across Multiple Processing Elements/Processors in Hardware," issued on December 30, 2014. XtreamEdge is the current assignee and owner of

the ’596 Patent and Concurrent Ventures has exclusively licensed all substantial rights to the ’596 Patent. Plaintiffs own all necessary rights, title, and interest in the ’596 Patent to bring this action, including the right to seek damages, including past damages, for any infringement thereof. A true and correct copy of the ’596 Patent is attached hereto as Exhibit 1. The ’596 Patent claims patent-eligible subject matter and is valid and enforceable.

13. U.S. Patent No. 10,873,753 (“the ’753 Patent”), titled “Hardware Defined Anything In A Platform With Swappable Pods, Message Interface, Sandboxes And Memory Superposition,” issued on December 22, 2020. XtreamEdge is the current assignee and owner of the ’753 Patent and Concurrent Ventures has exclusively licensed all substantial rights to the ’753 Patent. Plaintiffs own all necessary rights, title, and interest in the ’753 Patent to bring this action, including the right to seek damages, including past damages, for any infringement thereof. A true and correct copy of the ’753 Patent is attached hereto as Exhibit 2. The ’753 Patent claims patent-eligible subject matter and is valid and enforceable.

14. U.S. Patent No. 10,985,943 (“the ’943 Patent), titled “Hardware Defined Anything In A Platform With Swappable Pods, Message Interface, Sandboxes And Memory Superposition,” issued on April 20, 2021. XtreamEdge is the current assignee and owner of the ’943 Patent and Concurrent Ventures has exclusively licensed all substantial rights to the ’943 Patent. Plaintiffs own all necessary rights, title, and interest in the ’943 Patent to bring this action, including the right to seek damages, including past damages, for any infringement thereof. A true and correct copy of the ’943 Patent is attached hereto as Exhibit 3. The ’943 Patent claims patent-eligible subject matter and is valid and enforceable.

15. U.S. Patent No. 10,944,634 (“the ’634 Patent”), titled “Optimization for Network Connections,” issued on March 9, 2021. XtreamEdge is the current assignee and owner of the

'634 Patent and Concurrent Ventures has exclusively licensed all substantial rights to the '634 Patent. Plaintiffs own all necessary rights, title, and interest in the '634 Patent to bring this action, including the right to seek damages, including past damages, for any infringement thereof. A true and correct copy of the '634 Patent is attached hereto as Exhibit 4. The '634 Patent claims patent-eligible subject matter and is valid and enforceable.

16. U.S. Patent No. 9,529,767 ("the '767 Patent"), titled "System And Method For Abstracting SATA And/Or SAS Storage Media Devices Via A Full Duplex Queued Command Interface To Increase Performance, Lower Host Overhead, And Simplify Scaling Storage Media Devices And Systems," issued on December 27, 2016. XtreamEdge is the current assignee and owner of title of the '767 Patent and Concurrent Ventures has exclusively licensed all substantial rights to the '767 Patent. Plaintiffs own all necessary rights, title, and interest in the '767 Patent to bring this action, including the right to seek damages, including past damages, for any infringement thereof. A true and correct copy of the '767 Patent is attached hereto as Exhibit 5. The '767 Patent claims patent-eligible subject matter and is valid and enforceable.

IV.
JURISDICTION AND VENUE

17. Plaintiffs incorporate by reference paragraphs 1-16 herein.

18. This civil action arises under the patent laws of the United States, 35 U.S.C. § 1 *et seq.*, including without limitation 35 U.S.C. §§ 271, 281, 283, 284, and 285. This is a patent infringement lawsuit over which this Court has subject matter jurisdiction under, *inter alia*, 28 U.S.C. §§ 1331 and 1338(a).

19. This District has general and specific personal jurisdiction over Defendants because Defendants have committed acts within this District giving rise to this action, including, on information and belief, developing, making, using, marketing, selling, and testing the infringing

products; transact and conduct business in this District and the State of Texas; and transact and conduct business with residents of this District and the State of Texas.

20. Plaintiffs' causes of action arise, at least in part, from Defendants' contacts with and activities in and/or directed at this District and the State of Texas. Defendants have systematic and continuous business activities in this District. As described below, Defendants have committed acts of patent infringement giving rise to this action within this District.

21. Defendants have infringed and continue to infringe the Asserted Patents within this District and the State of Texas by making, using, selling, licensing, offering for sale, and/or importing or exporting in, into, or out of this District and elsewhere in the State of Texas, products and services covered by claims in the Asserted Patents, including without limitation products that, when made or used, practice the claimed methods and systems of the Asserted Patents. Defendants, directly and through intermediaries, make, use, sell, offer for sale, import, ship, distribute, advertise, promote, and/or otherwise commercialize such infringing products and services in or into this District and the State of Texas. Defendants regularly conduct and solicit business in, engage in other persistent courses of conduct in, and/or derive substantial revenue from goods and services provided to residents of this District and the State of Texas.

22. This Court has personal jurisdiction over Defendants pursuant to TEX. CIV. PRAC. & REM. CODE § 17.041 *et seq.*, the Texas Long Arm Statute.

23. Venue is proper in this District under 28 U.S.C. §§ 1391 and 1400(b).

24. Defendants are doing business, either directly or through respective agents, on an ongoing basis in this District and elsewhere in the United States and have committed acts of infringement in this District. AMD and, on information and belief Pensando, have a regular and established place of business in this Judicial District, including at their Austin campus, located at

7171 Southwest Parkway, Austin, TX 78735. On information and belief, Defendants make, use, sell, test, offer to sell, and/or import infringing products into and/or within this District, including at its Austin campus. Defendants maintain a permanent and/or continuing presence within this District at their Austin campus, and have the requisite minimum contacts with this District such that this venue is a fair and reasonable one. Upon information and belief, Defendants have transacted and, at the time of the filing of the Complaint, are continuing to transact business within this District.

25. Defendants purposefully direct or control the sale of the Accused Products, including, on information and belief, for sale in Texas and elsewhere in the United States, and expect and intend that the Accused Products will be so sold in this District. Defendants purposefully place the Accused Products—whether by itself or through subsidiaries, affiliates, or third parties—into an international supply chain, knowing that the Accused Products will be sold in the United States, including Texas and in this District. Therefore, Defendants also facilitate the sale of the Accused Products in Texas. On information and belief, Defendants know that Texas is a termination point of its established distribution channels.

26. On information and belief, current and former Pensando and AMD employees with knowledge relevant to the claims in this case are located in this District and within the State of Texas. For all of these reasons, venue is appropriate and convenient in this District.

V.
DEFENDANTS ARE INTENTIONALLY AND KNOWINGLY USING PLAINTIFFS' TECHNOLOGY

27. Plaintiffs incorporate by reference paragraphs 1-26 herein.

28. On information and belief, Defendants have infringed and continue to infringe one or more claims of the Asserted Patents by, at a minimum, making, using, offering for sale, and selling infringing products and services in the United States and in this District.

29. The Accused Products are all products including or related to AMD's DPU technology, including without limitation, AMD's Pensando Data Processing Units including the AMD Pensando Giglio Data Processing Unit, the Elba Data Processing Unit, the Capri Data Processing Unit, and future versions, such as the announced Salina Data Processing Unit, the AMD Pensando Distributed Services Card, the AMD Pensando Distributed Services Platform, the AMD Pensando DPU system, the AMD Pensando SmartNIC, the AMD Pensando SmartSwitch, the AMD Pensando Software-In-Silicon Development Kit ("SSDK") and reference pipelines, as well as any products incorporating those items (collectively, the "Accused Products/Services"). Each of these products and services practice at least one claim of each Asserted Patent.

30. In 2014, AMD held detailed technical discussions with Concurrent Ventures' commercialization agent, HellaStorm, Inc., during which AMD was provided with the specifics of Concurrent Ventures' technology. HellaStorm and AMD entered into a "Mutual Nondisclosure Agreement" for the purpose of "Internal evaluation and/or testing," which was signed by Jesse Beeson, the CEO of Concurrent Ventures, and Harry Wolin, AMD's Senior Vice President and General Counsel. The parties met on-site at AMD in June 2014, and the meeting included a number of AMD executives, including Anil Rao, the then Corporate Vice President of Products who "[d]rove data center platforms, systems and software for AMD's Server Business Unit"[4] and

[4] https://www.linkedin.com/in/anilrrao.

Dhiraj Mallick, the then General Manager and Corporate Vice President responsible "for $20B data center group" and "next-generation data center architectures, solutions, and innovations."[5]

31. After that meeting, AMD was sent a proposed engagement plan, outlining a potential project. The parties did not ultimately agree to this proposed project. On information and belief, Defendants have been on notice of the Asserted Patents and their infringement since at least these discussions and, for later issuing patents, since their issuance, as evidenced by the above facts; the similarities between Defendants' Accused Products and their marketing and the Plaintiffs' patented technology and marketing; and the small nature of the industry, Plaintiffs' pioneering role in that market, and Defendants' direct competition against Plaintiffs. Defendants have also been on notice of the Asserted Patents and its specific infringement since the filing of this complaint.

32. Defendants' infringement of the Asserted Patents has been and continues to be willful in view of the above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

VI.
FIRST CLAIM
Infringement of the '596 Patent

33. Plaintiffs incorporate by reference paragraphs 1-32 herein.

34. On information and belief, Defendants have infringed and continue to infringe the '596 Patent in violation of 35 U.S.C. § 271(a), either literally or through the doctrine of equivalents, by making, using, selling, offering for sale, and/or importing into the United States products and/or methods that practice at least claim 1 of the '596 Patent, including the Accused Products.

[5] https://www.linkedin.com/in/dhirajmallick/details/experience.

35. Each of the Accused Products meet every limitation of claim 1 of the '596 Patent, which recites:

> 1. A system for dividing and synchronizing a processing task across a plurality of processing elements comprising:
>
> an input queue implemented in hardware;
>
> an output queue implemented in hardware;
>
> a first processing element having access to said input queue and said output queue;
>
> at least one second processing element in communication with said first processing element;
>
> a reservation register implemented in hardware storing a value indicative of available space in said input queue, said reservation register accessible by both said first processing element and said at least one second processing element;
>
> computer storage storing instructions, which when executed by said at least one second processing element:
>
> > accesses said reservation register and reads said stored value;
> >
> > determines when said read value indicates available space in said input queue for said first processing element to issue a command;
> >
> > notifies said first processing element to issue said command to said input queue; and
>
> wherein said first processing element receives notification from said at least one second processing element regarding issuing the command, issues said command to said input queue, and receives a response corresponding to said command from said output queue.

36. The Accused Products, including AMD Pensando's DPU technology and DPU-enabled systems, allow for hardware processing of specific functions on the data path with hardware queues supporting various operations implemented in hardware. The Accused Products include systems for dividing and synchronizing a processing task across a plurality of processing

elements. For example, the AMD Pensando DPUs "offload infrastructure services from the computing CPU and make use of hardware accelerators to boost performance of those services" and "improve security by running security softw[a]re like distributed firewalls (on the DPU) on different cores than the workloads (on the x86 CPU)." https://www.delltechnologies.com/asset/en-us/products/networking/briefs-summaries/smartdpu-software-solutions-solution-brief.pdf; *see* https://www.amd.com/system/files/documents/pensando-dsc-200-product-brief.pdf; https://www.servethehome.com/pensando-distributed-services-architecture-smartnic; https://pensando.io/project-monterey-early-access-program (accessed May 6, 2023).

37. The Accused Products include an input queue implemented in hardware and an output queue implemented in hardware. For example, the AMD Pensando Products include numerous hardware queues, such as host or ARM processes, which initiate multi-level scheduler to inject P4 pipeline tokens. *See* https://www.servethehome.com/pensando-distributed-services-architecture-smartnic; *see also* https://hc32.hotchips.org/assets/program/conference/day2/HotChips2020_Networking_Pensando_v3.pdf. For another example, the Accused DPU products contain input queues (for example, P4 RxDMA) and output queues (for example, P4 TxDMA). *Id.*

38. The Accused Products include a first processing element having access to said input queue and said output queue, for example a CPU on a host system having access to said input queue and said output queue. As examples, hosts within the system and Virtual Machines ("VMs") running on the system containing a DPU, as well as processors within the DPU, have access to the input queue for (as an example) scheduling operations performed by parts of the DPU. *See* https://www.youtube.com/watch?v=in7athW-PaQ at 17:11.

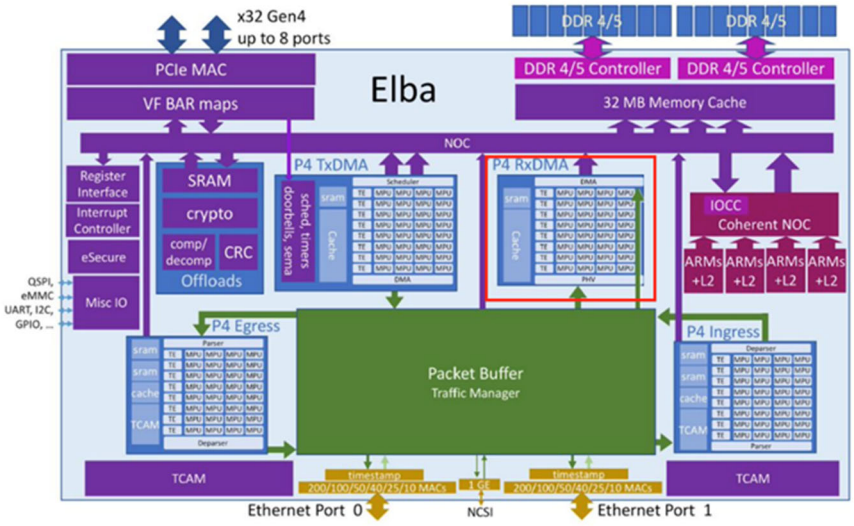

https://www.servethehome.com/pensando-distributed-services-architecture-smartnic.

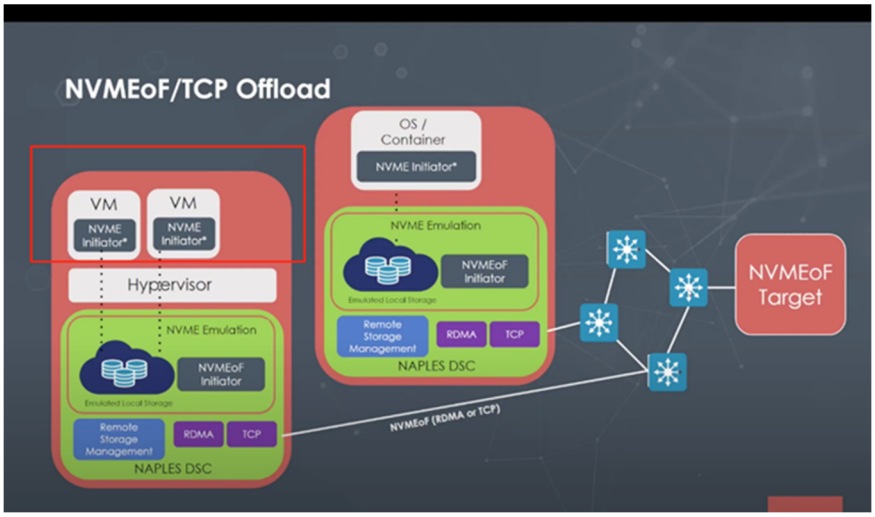

https://www.youtube.com/watch?v=in7athW-PaQ at 17:11.

39. The Accused Products include at least one second processing element in communication with said first processing element, for example a second processing element connected via internal interconnects and a PCIe interface to the first processing element. *See* https://www.youtube.com/watch?v=in7athW-PaQ at 18:02.

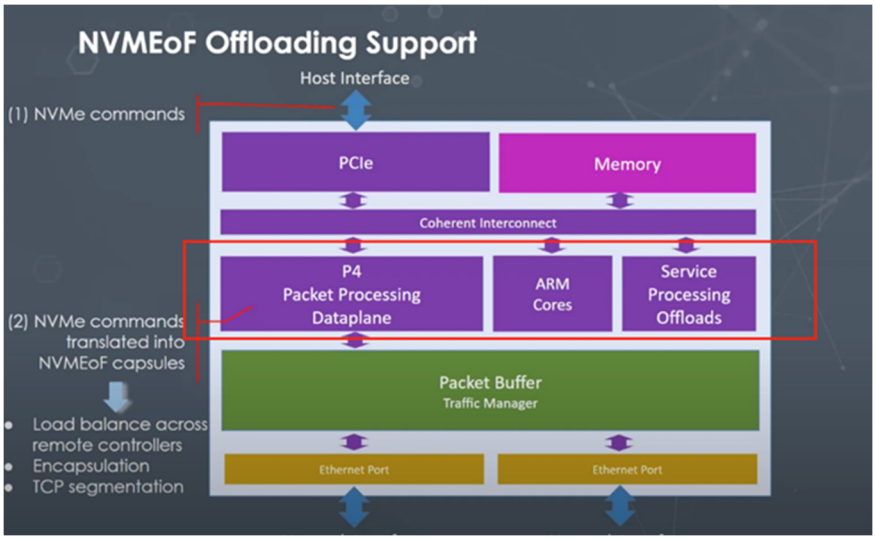

https://www.youtube.com/watch?v=in7athW-PaQ at 18:02.

40. The Accused Products contain a reservation register implemented in hardware storing a value indicative of available space in said input queue, said reservation register accessible by both said first processing element and said at least one second processing element accessible by the first and second processing element. *See* https://www.servethehome.com/pensando-distributed-services-architecture-smartnic. For example, the doorbells for the hardware queues initiate the scheduler to inject P4 pipeline tokens and this doorbell mechanism allows objects to be added to any queue. *Id.*

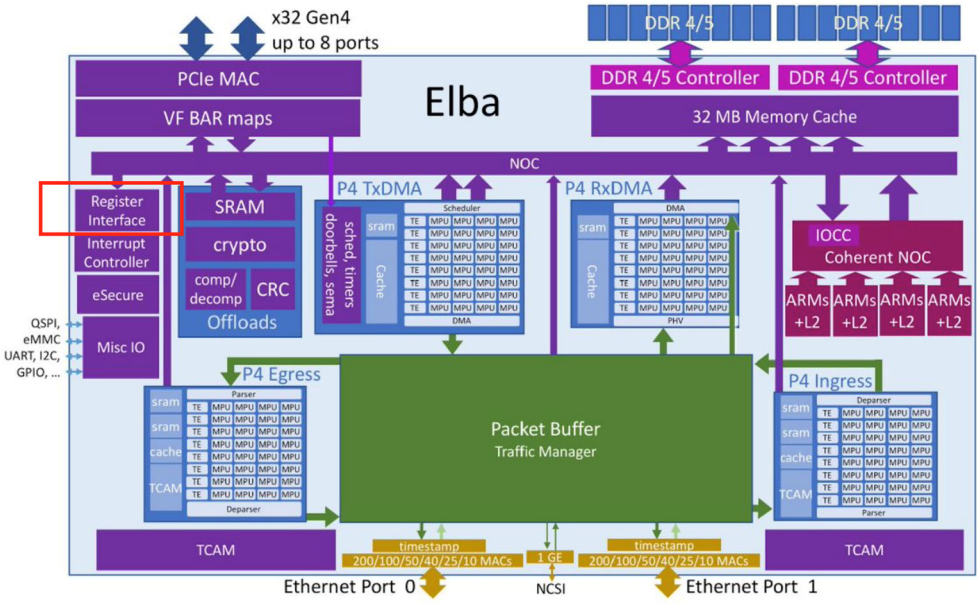

SOC, NOC, & hardware queues

- Network On Chip (NOC) connects CPUs, memory, P4 DMA, and offload engines in a cache-coherent domain
- ARM A-72 CPUs run SMP Linux, DPDK supported
- P4 programs can push packet headers and data structures directly into ARM L2 caches or the last level system cache
- Doorbells for 16 million hardware queues are mapped to host or ARM processes, initiates multi-level scheduler to inject P4 pipeline tokens

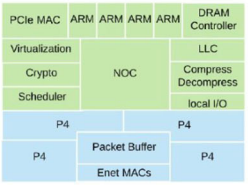

Hot Chips 32 Pensando SOC NOC And Hardware Queues

https://www.servethehome.com/pensando-distributed-services-architecture-smartnic.

New Constructs Required

- Doorbell
  - A register the driver can update to start a P4 program
  - A register a P4 program can update to notify completion
- DMA
  - GSO
    - Copy header and segments to form multiple packets
  - LRO
    - Copy payload from multiple packets to form a larger message
- Timers
- Stateful Memory
  - To keep track of state
  - GSO : last segment offloaded; LRO : last segment received

PENSANDO

https://opennetworking.org/wp-content/uploads/2020/04/Plenary-4-Slide-Deck.pdf.

Hardware Queue Management

Packets which enter the P4 pipeline form the wire or which originate from internal events are placed in hardware queues. Hardware queues are used to manage interfaces, flows, connections, and any other objects which require ordered tracking and scheduling. Software can configure up to 16 million hardware queues (Figure 5). Each queue stores its current state in a DRAM-based qstate record, which includes a count of enqueued objects, pointers to arrays or linked lists of enqueued objects, connection state and peer information, and a process ID used for memory protection.

http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf, p. 46.

Key Features

Feature	Description
Network Interfaces	56 Gbps PAM4 SerDes supporting Dual 200 GE, Quad-100/50/25/10GE 1GE management port
PCIe Interface	32 Lanes of PCIe Ge n4, configurable as root complex or end-point mode 2x16 / 4x8 / 8x4 with QOS support in multi-host applications
Data Pipeline	P4 data pipeline comprising 144 match processing units (MPUs) @2 GHz Provides high performance capabilities in packet and message processing, at line rate
SOC	16x Arm A72 CPU cores @3 GHz with I-Cache, D-Cache and LLC Cache QSPI Flash, EMMC storage for embedded OS Secure Boot with hardware root of trust
Memory	Dual DDR4-3200 interfaces supporting 8 GB - 64 GB system memory
Offloads	Inline IPsec and DTLS, Bulk Crypto, PKE, Compression, Decompression, Checksums, Deduplication, Erasure Coding
Scale	2K VNICs, 16 M hardware queues, highly scalable P4 tables (stateful and stateless) accessible at every stage of the pipeline
Scheduling	Queue Group scheduling with Min/Max rate High Priority Option per queue group P4 meters and QoS priorities
ROCEv2	Memory based scatter/gather Lists (SGL) in DMA commands Latency optimized hardware data path with all context on DPU
Single Wire Management	Connects 1GE BMC controller and Uplink ports in Standby Power mode

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf; https://www.amd.com/system/files/documents/pensando-giglio-product-brief.pdf.

41. The Accused Products contain computer storage storing instructions, which when executed by said at least one second processing element, accesses the reservation register and reads the stored value using the scheduler. https://www.amd.com/system/files/documents/pensando-dsc-200-product-brief.pdf; http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf. The Accused Products determine when said read value indicates available space in said input queue for said first processing element to issue a command. https://www.servethehome.com/pensando-distributed-services-architecture-smartnic; http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf. With the Accused Products, the second processing element notifies said first processing element to issue

said command to said input queue. https://www.youtube.com/watch?v=in7athW-PaQ; https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf. For example, once the DPU determines available space, the first processing element issues a command utilizing the scheduler and the reservation register.

42. Ultimately, the first processing element receives notification from said at least one second processing element regarding issuing the command, issues said command (via the host interface) to said input queue and receives a response corresponding to said command from said output queue. https://www.servethehome.com/pensando-distributed-services-architecture-smartnic; https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf; http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf; https://opennetworking.org/wp-content/uploads/2020/04/Mario-Baldi-Slide-Deck.pdf.

43. On information and belief, Defendants' infringement through its use of its DPU technology and DPU-enabled systems, described above, is exemplary of its infringement with respect to all the Accused Products.

44. Defendants have also indirectly infringed, and continue to indirectly infringe, the '596 Patent under 35 U.S.C. § 271(b) and (c).

45. Defendants knowingly and intentionally actively aided, abetted, and induced others to directly infringe at least claim 1 of the '596 Patent (such as its customers in this District and throughout the United States), and continue to do so, by, for example, selling and offering access to and encouraging and supporting use of the Accused Products.

46. Defendants contributed to the direct infringement of at least claim 1 of the '596 Patent under 35 U.S.C. § 271(c), and continue to do so, by, for example, supplying, with knowledge of the '596 Patent, a material part of a claimed invention, where the material part is not a staple

article of commerce and is incapable of substantial noninfringing use. For example, Defendants have provided, owned, operated, sold, offered to sell, leased, licensed, used, and/or imported, and continue to do so, various hardware and/or software that make up and enable the Accused Products, including as used in third-party (including customer) systems (including as discussed above), are a material part of the claimed invention, are not a staple article of commerce, and are incapable of substantial non-infringing uses.

47. As explained above, Defendants' infringement has been and continues to be willful in view of the facts asserted above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

VII.
SECOND CLAIM
(Infringement of the '753 Patent)

48. Plaintiffs incorporate by reference paragraphs 1-47 herein.

49. On information and belief, Defendants have infringed and continue to infringe the '753 Patent in violation of 35 U.S.C. § 271(a), either literally or through the doctrine of equivalents, by making, using, selling, offering for sale, and/or importing into the United States products and/or methods that practice at least claim 1 of the '753 Patent, including the Accused Products.

50. Each of the Accused Products meet every limitation of claim 1 of the '753 Patent, which recites:

> 1. A platform for data flow processing, comprising:
> one or more swappable pods or cards in one or more chassis, coupled through a messaging interface network;
>
> each of the one or more swappable pods or cards having one or more hardware modules or one or more software modules;

one or more of the plurality of swappable pods or cards having a portion for user-definable hardware modules or user-definable software modules;

the plurality of swappable pods or cards being user-configurable to implement data flow processing architectures; and

a network coupled to the one or more swappable pods or cards and supporting messaging-based communication using packets each having a header with a chassis identifier, a board identifier, a module identifier, an instance identifier, and a type identifier, so that each type of module, each instance of a type of module, and each module on each board in each chassis can be addressed through the header.

51. The Accused Products include a platform for data flow processing comprising one or more swappable pods or cards in one or more chassis coupled through a messaging interface network by, for example, supporting the P4 programming language.

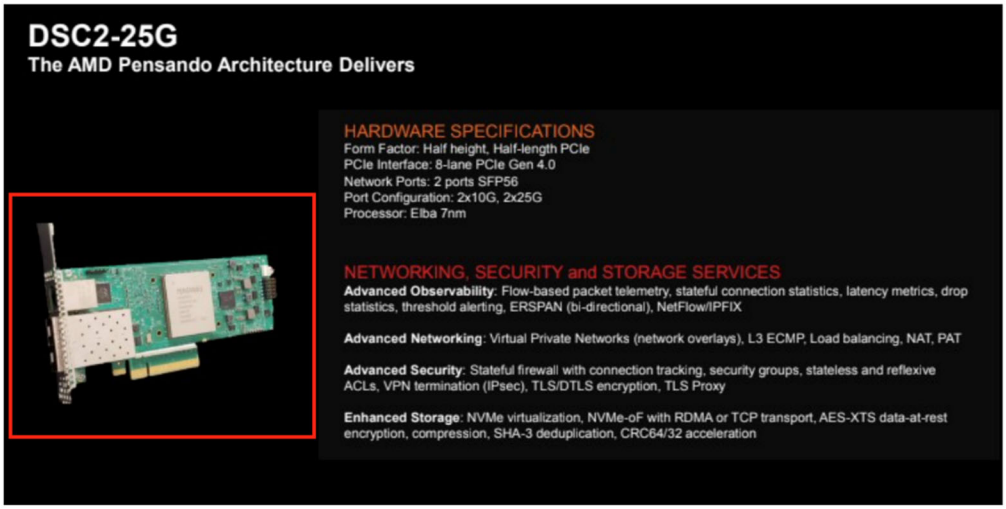

AMD Pensando DSC2 25G

https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc; *see* http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf, p. 44 (stating that "Pensando ASICs include multiple P4 pipelines."); https://opennetworking.org/news-and-events/blog/pensando-announces-p4-programmable-platform-and-joins-p4-community.

52. The Accused Products include swappable pods or cards in, for example, the Pensando DSC which has hardware modules and software modules and, for example, the Pensando DSC contains user-definable hardware (via P4) modules and user-definable software modules (via ARM cores). https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-

ipu-nvidia-bluefield-marvell-octeon-fungible;

https://www.amd.com/system/files/documents/pensando-dsc-200-product-brief.pdf.

53. The Accused Products are designed to have portions that are user-configurable to implement data flow processing architectures such as for example a SmartSwitch. https://www.theregister.com/2022/04/07/amd_pensando_aws; https://www.networkworld.com/article/3690334/aruba-to-prioritize-sase-private-5g-data-center-networking.html.

54. The Accused Products include a network coupled to the one or more swappable pods or cards.

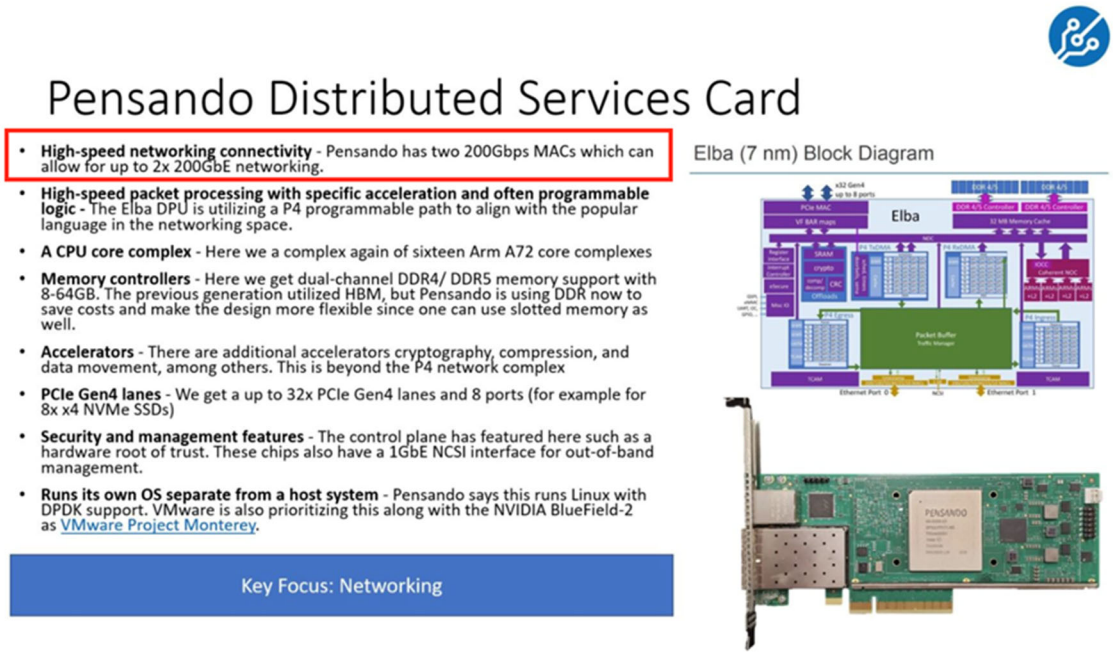

Pensando DSC DPU Example Q2 2021

https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible.

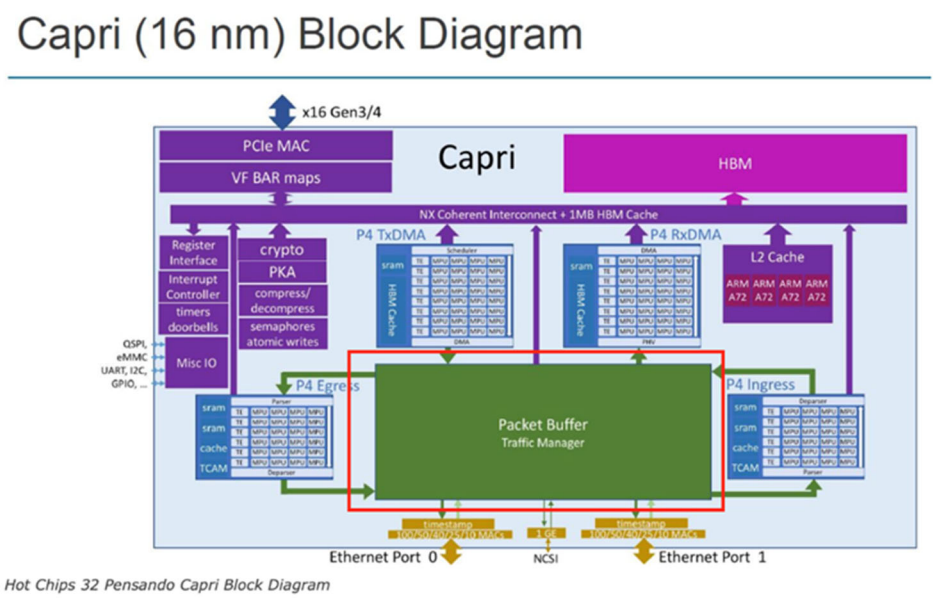

https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible; *see* https://www.amd.com/system/files/documents/pensando-smartswitches.pdf.

55. The Accused Products support messaging-based communication using packets each having a header with a chassis identifier, a board identifier, a module identifier, an instance identifier, and a type identifier, so that each type of module, each instance of a type of module, and each module on each board in each chassis can be addressed through the header. For example, specifically, the Pensando DPUs are programmable using P4 which allows for the customization of how a network device processes packets. https://packetpushers.net/aruba-puts-dpus-into-new-top-of-rack-switch-5-questions; https://www.amd.com/system/files/documents/pensando-smartswitches.pdf; https://www.amd.com/system/files/documents/pensando-project-monterey.pdf.

The Key Difference in Architecture: The Pensando Flow-Based Engine

We have previously mentioned that:

SmartSwitches have a different forwarding mechanism that is stateful. Each packet between host A and host B is processed as a part of a session between A and B. The session comprises two unidirectional flows, one from A to B and one from B to A.

Sessions and flows are identified using an extended five-tuple that includes the basic five fields (IP source and destination addresses, protocol type, TCP/UDP source destination ports) and other information like input interface, MAC addresses, TCP flags, etc.

When you build hardware capable of implementing this new forwarding mechanism at wire speed, many new features become possible.

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf.

multicast or span replication; then enter the P4 egress pipeline to apply, for example, network address translation, policers, and telemetry. Having completed network processing tasks, the packet can be sent back out an Ethernet MAC port or it can be passed up toward the ARM subsystem or local host.

If the packet is sent to the embedded ARM or local host, it must be processed according to the needs of the kernel or user level driver running on that CPU. The packet passes through a third P4 pipeline, where it runs another program in accordance with the protocol type of the packet, i.e., Ethernet, RDMA, NVMe, or a new protocol. A significant advantage of running protocol-specific P4 programs associated with a targeted kernel or user driver is the ability to customize descriptor formats, completion events, interrupt schemes, and data formats to the needs of the targeted driver.

In summary, network packets enter the ASIC, pass through P4 ingress, and P4 egress pipelines to apply security and network services, are associated with a target interface and device type, and pass through P4 driver interface programs. The final packet or data buffer are delivered to the attached host, embedded ARM, or integrated offload engine. Packets sent by the host enter through the PCIe MAC to be processed by the P4 driver interface programs, then by the P4 ingress and egress pipelines to apply security and network services before being sent out through the Ethernet MAC. The following section takes a closer look at the P4 pipeline internals.

P4 Pipelines

Pensando ASICs include multiple P4[3] pipelines (Figure 2) as the centerpiece of a domain-specific processing architecture. Each P4 pipeline starts with a programmable parser. At its core, the programmable parser

FIGURE 2. Single P4 pipeline with 8 stages.

includes multiple engines designed to extract header fields and populate a packet header vector (PHV) at rates of 100 million packets per second or more. Following the parser are six to eight match action stages, each stage able to load multiple tables per packet and perform complex actions which modify the PHV and update memory data structures. A single P4 pipeline includes local TCAM and SRAM resources where high bandwidth tables can be stored. Large scale tables as well as tables shared between multiple pipelines are stored in attached DRAM or host memory. The table base address steers table requests to SRAM or to DRAM, with per-pipeline caches reducing DRAM traffic for commonly accessed table entries. The P4 pipeline concludes with a deparser to reassemble the packet from the updated PHV fields. Alternatively, the pipeline can conclude with a DMA engine to directly copy packet header, payload, and meta data structures to and from local or host memory for CPU or offload engine processing.

P4 Stages

Each P4 stage (Figure 3) within a pipeline begins with a table engine (TE). The TE extracts bits and bytes from the PHV in any combination to build a table key. Table keys constructed from multiple header fields can be up to 512 bits in width, or keys can be chained together to create extra wide keys up to 2048 bits in width. The resulting key may be hashed or used as a direct index, then matched against any data structure in TCAM, SRAM, DRAM, or attached host memory.

The table data and lookup result (match or no match) are then forwarded to a match processing unit

http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf, p. 44.

56. On information and belief, Defendants' infringement through its use of its DPU technology and DPU-enabled systems, described above, is exemplary of its infringement with respect to all the Accused Products.

57. Defendants have also indirectly infringed, and continue to indirectly infringe, the '753 Patent under 35 U.S.C. § 271(b) and (c).

58. Defendants knowingly and intentionally actively aided, abetted, and induced others to directly infringe at least claim 1 of the '753 Patent (such as its customers in this District and throughout the United States), and continue to do so, by, for example, selling and offering access to and encouraging and supporting use of the Accused Products.

59. Defendants contributed to the direct infringement of at least claim 1 of the '753 Patent under 35 U.S.C. § 271(c), and continue to do so, by, for example, supplying, with knowledge of the '753 Patent, a material part of a claimed invention, where the material part is not a staple article of commerce and is incapable of substantial noninfringing use. For example, Defendants have provided, owned, operated, sold, offered to sell, leased, licensed, used, and/or imported, and continue to do so, various hardware and/or software that make up and enable the Accused Products, including as used in third-party (including customer) systems (including as discussed above), are a material part of the claimed invention, are not a staple article of commerce, and are incapable of substantial non-infringing uses.

60. As explained above, Defendants' infringement has been and continues to be willful in view of the facts asserted above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

VIII.
THIRD CLAIM
(Infringement of the '943 Patent)

61. Plaintiffs incorporate by reference paragraphs 1-60 herein.

62. On information and belief, Defendants have infringed and continue to infringe the '943 Patent in violation of 35 U.S.C. § 271(a), either literally or through the doctrine of equivalents, by making, using, selling, offering for sale, and/or importing into the United States products and/or methods that practice at least claim 1 of the '943 Patent, including the Accused Products.

63. Each of the Accused Products meet every limitation of claim 1 of the '943 Patent, which recites:

> 1. A programmable device for data flow processing in a user-configurable server with swappable pods or cards, comprising:
>
> a programmable logic device (PLD) having a first region and a second region;
>
> the first region comprising a hardware-based or firmware-based router with a port enabling communication between the router and the second region; and the second region comprising one or more sandboxes with user-definable programmable electronic circuits of the PLD,
>
> wherein the port comprises: a bridge being lockable to prevent user access, and unlockable to enable user access, to each of the sandboxes and the user-definable programmable electronic circuits therein.

64. The Accused Products include a programmable device for data flow processing in a user-configurable server with swappable pods or cards, for example, supporting the P4 programming language.

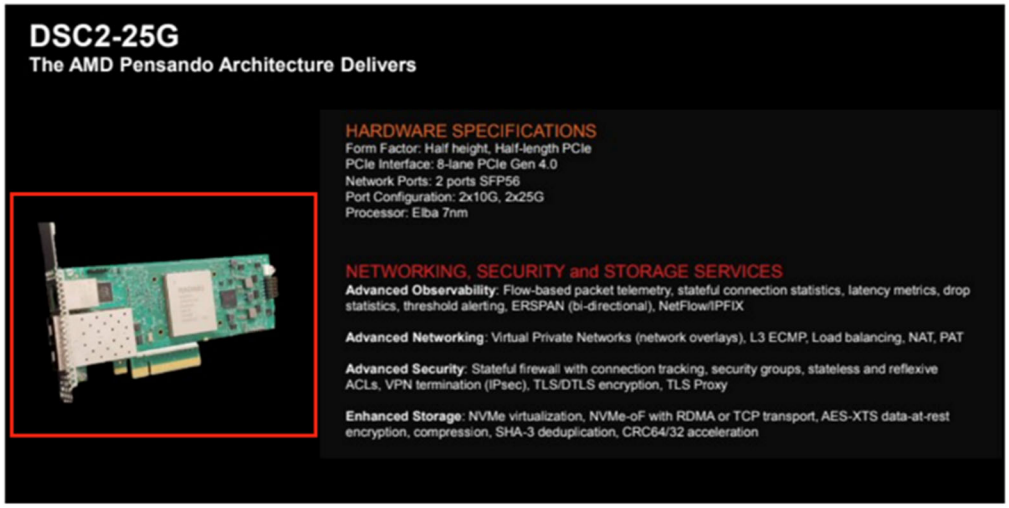

AMD Pensando DSC2 25G

https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc; http://gibsonnet.net/aix/ibm/mmi202102.issue.pdf (stating that "Pensando ASICs include multiple P4 pipelines.").

65. The Accused Products contain a programmable logic device ("PLD") having a first region and a second region, for example, the AMD DSC is a programmable device. https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible; https://www.amd.com/system/files/documents/aruba-cx-10000-series-datasheet.pdf; https://www.amd.com/en/solutions/infrastructure-acceleration. For example, the DPU within the Pensando DSC contains the following exemplary first and second regions.

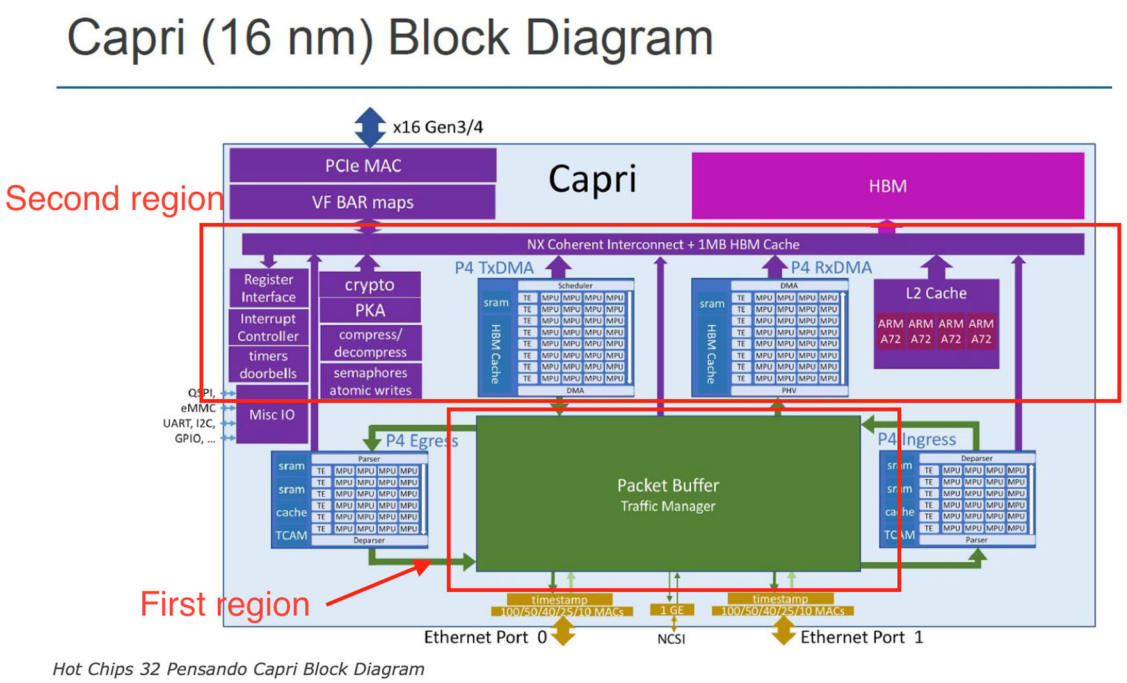

https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible.

66. The first region comprising a hardware-based or firmware-based router with a port enabling communication between the router and the second region. For example, the first region of the Accused Products includes a packet buffer traffic manager for routing data between the network and the second region.

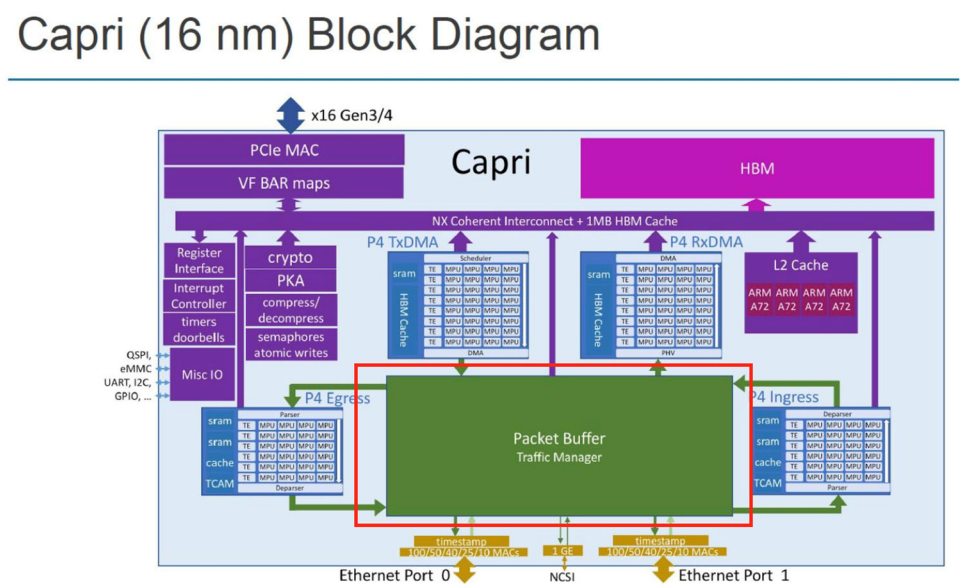

Hot Chips 32 Pensando Capri Block Diagram

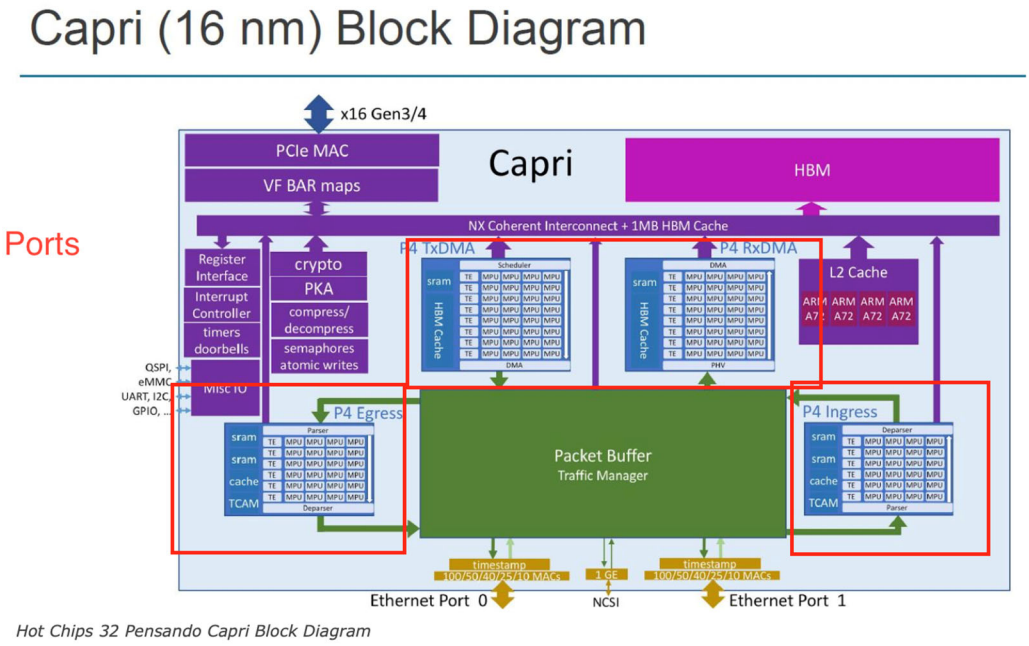

Hot Chips 32 Pensando Capri Block Diagram

https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible.

67. The second region comprises one or more sandboxes with user-definable programmable electronic circuits of the PLD. For example, the Accused Produces have one or more sandboxes with user-definable programmable electronic circuits of the PLD. For example, the second region of the Accused DSC contains the P4 programmable components and the ARM

CPU cores. https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible.

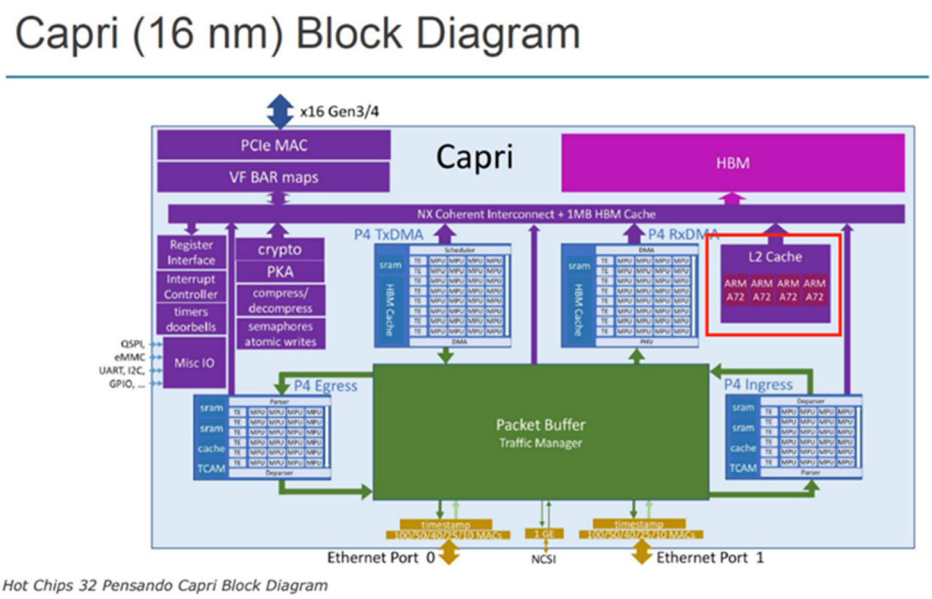

https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc.

68. The port comprises: a bridge being lockable to prevent user access, and unlockable to enable user access, to each of the sandboxes and the user-definable programmable electronic circuits therein. For example, the Accused DSC has a port that comprises: a bridge being lockable to prevent user access, and unlockable to enable user access, to each of the sandboxes and the user-definable programmable electronic circuits therein.

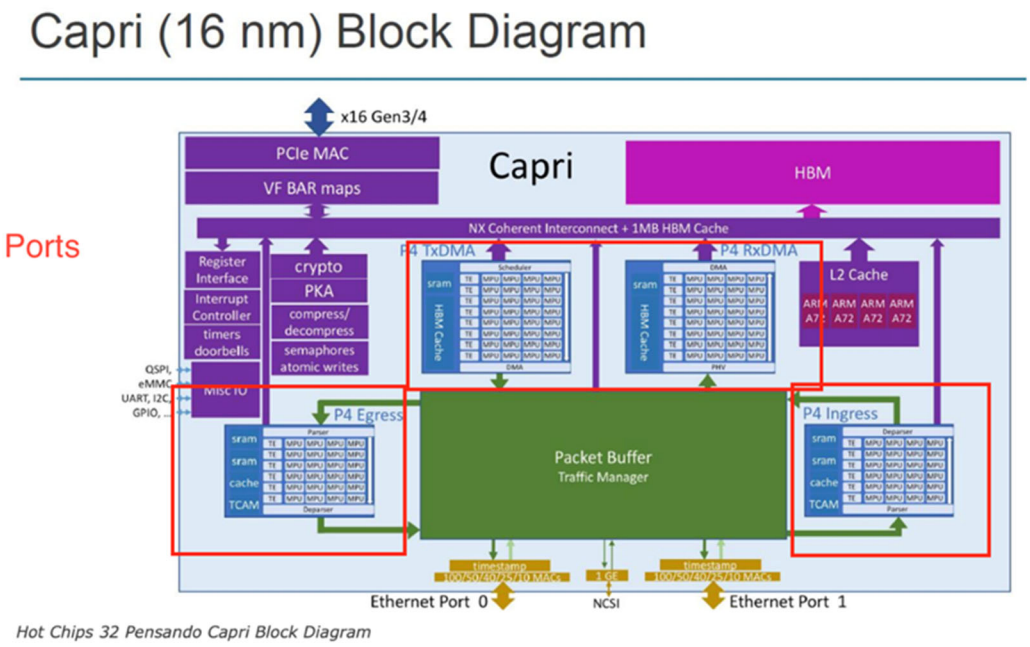

Hot Chips 32 Pensando Capri Block Diagram

https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible.

Encryption Services Architecture

The IPsec data path as well as handshake and key management elements are implemented within the DSC, assuring an extremely high level of security, as both keys and sensitive encryption policies are maintained inside the secure perimeter of the DSC ASIC. The PenAccel™ subsystem implements high-performance hardware-based crypto acceleration in the DSC data path to provide IPsec encryption services at line speed (25Gbps or 100Gbps). The on-chip ARM subsystem runs the crypto handshake and key management control plane for the associated crypto protocol (e.g. TLS handshake or IKE for IPsec).

The default ciphersuite for IPsec ESP operation is AES-GCM-256, providing very high security, CNSA-compliant encryption and payload integrity and authentication.

In order to support very high crypto connection rates, hardware acceleration of both public key and symmetric encryption algorithms is available to the ARM subsystem, including RSA, Elliptic Curve, and Diffie Hellman. In addition, a high-grade entropy-based random number generator, compliant with NIST SP 800-90A/B, is used.

Platform Security

The *PenTrust* subsystem is the Root of Trust (RoT) for the Pensando Programmable Services Processor. PenTrust is the first subsystem to boot at power-on or hardware reset from an immutable embedded ROM, making it the first link in the Secure Boot Chain of Trust. The firmware for the PenTrust subsystem is burned in ROM during ASIC manufacturing and cannot be modified or tampered with once it is programmed. Together with a chip-unique key and identity, the PenTrust subsystem validates all code that runs on the DSC with a digital signature.

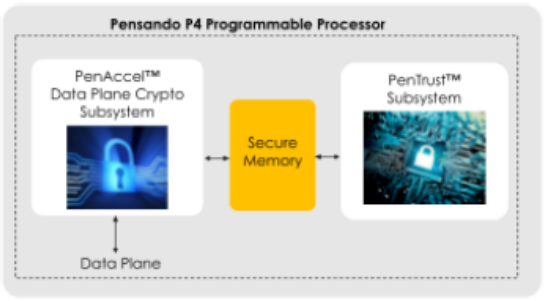

Ex. 6, Pensando IPsec Solutions, Pensando White Paper (2021); *see* https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf.

Segmentation is key to preventing unwanted lateral movement, by inspecting all East-West traffic in the data center and applying policies that stop bad actors from moving through the internal network. Several approaches have been tried to achieve segmentation in the past, but with very limited success:

- Hardware **next-generation firewall** appliance-based segmentation
- **Virtualized firewall appliance**-based segmentation
- **Software agent**-based segmentation
- **Network switch** (stateless ACL)-based segmentation

https://www.amd.com/system/files/documents/zero-trust-building-secure-data-centers-cx-10000.pdf.

PenTrust Subsystem

Internal to Pensando's Capri silicon at the heart of each Distributed Services Card (DSC) is a security element called the PenTrust Subsystem that is responsible for the fundamental protection mechanisms in the device. The module is isolated from the rest of the chip and has its own CPU, ROM, RAM, and cryptographic engines, similar to Trusted Platform Modules (TPM) used in server systems. *PenTrust* accesses chip resources outside of its secure perimeter via its own bus mastering DMA engine, and it can receive requests (e.g. to create keys or sign certificates) via a narrow, secure-access interface. Other modules can deposit requests and asynchronously pick up responses as they become available.

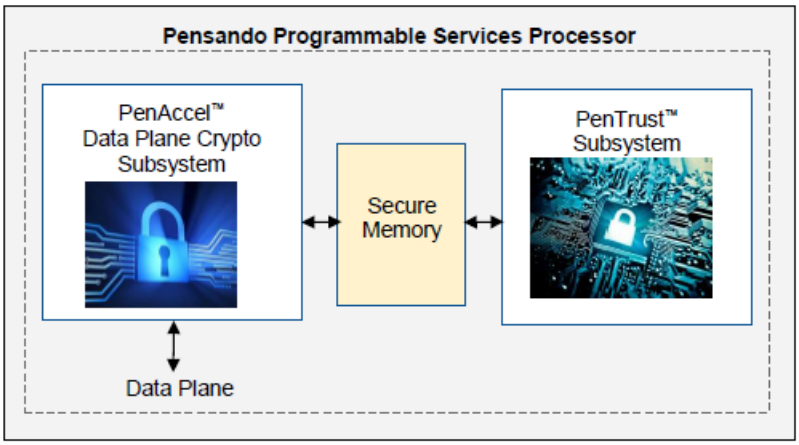

Fig 2. Capri P4 Programmable Processor

The *PenTrust* subsystem provides trusted crypto services such as asymmetric (public key) sign and verify operations, as well as symmetric (secret key) encryption/decryption and secure hashing operations. It also enables secure key store and management via primitives to generate keys, wrap/unwrap keys, and import/export them in encrypted form.

Root of Trust

The *PenTrust* subsystem is the Root-of-Trust (RoT) for the Pensando Programmable Services Processor. As discussed in the Secure Boot section below, *PenTrust* is the first subsystem to boot at power-on, and it does so from its immutable embedded ROM; making it the first link in the secure boot chain. At the heart of the RoT is a Physically Unclonable Function (PUF). This is a specialized silicon element, tied to specific physical properties of each silicon die, that provides a device-specific seed key. The result is that each chip produced by Pensando will have a unique PUF seed key that cannot be read outside of the device nor tampered with. The seed key is exclusively used to derive a 256-bit AES Storage Root Key (SRK) and an ECDSA-P384 Endorsement Key (EK).

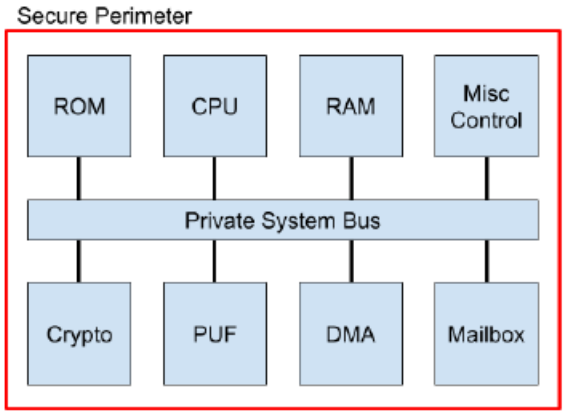

Fig 3. PenTrust Subsystem

Ex. 7, Security Architecture of Pensando's Distributed Services Platform, Pensando White Paper (2020).

69. On information and belief, Defendants' infringement through its use of its DPU technology and DPU-enabled systems, described above, is exemplary of its infringement with respect to all the Accused Products.

70. Defendants have also indirectly infringed, and continue to indirectly infringe, the '943 Patent under 35 U.S.C. § 271(b) and (c).

71. Defendants knowingly and intentionally actively aided, abetted, and induced others to directly infringe at least claim 1 of the '943 Patent (such as its customers in this District and throughout the United States), and continue to do so, by, for example, selling and offering access to and encouraging and supporting use of the Accused Products.

72. Defendants contributed to the direct infringement of at least claim 1 of the '943 Patent under 35 U.S.C. § 271(c), and continue to do so, by, for example, supplying, with knowledge of the '943 Patent, a material part of a claimed invention, where the material part is not a staple article of commerce and is incapable of substantial noninfringing use. For example, Defendants have provided, owned, operated, sold, offered to sell, leased, licensed, used, and/or imported, and continue to do so, various hardware and/or software that make up and enable the Accused Products, including as used in third-party (including customer) systems (including as discussed above), are a material part of the claimed invention, are not a staple article of commerce, and are incapable of substantial non-infringing uses.

73. As explained above, Defendants' infringement has been and continues to be willful in view of the facts asserted above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

IX.
FOURTH CLAIM
(Infringement of the '634 Patent)

74. Plaintiffs incorporate by reference paragraphs 1-73 herein.

75. On information and belief, Defendants have infringed and continue to infringe the ’634 Patent in violation of 35 U.S.C. § 271(a), either literally or through the doctrine of equivalents, by making, using, selling, offering for sale, and/or importing into the United States products and/or methods that practice at least claim 8 of the ’634 Patent, including with the Accused Products.

76. Each of the Accused Products meet every limitation of claim 8 of the ’634 Patent, which recites:

> 8. A tangible, non-transitory, computer-readable media having instructions thereupon which, when executed by a processor, cause the processor to perform a method comprising:
>
> collecting, at a first endpoint device coupled to a network, parameter values for determination of bandwidth of network connections to further endpoint devices;
>
> determining at a tuner server, distinct from the first endpoint device and the further endpoint devices, that a next network connection from the first endpoint device to a second endpoint device matches a geographical area of a past network connection; and
>
> initiating the next network connection, from the first endpoint device to the second endpoint device, based on the determination at the tuner server, with a transmission bandwidth based on the parameter values for the past network connection.

77. The Accused Products are tangible, non-transitory, computer-readable media having instructions thereupon which, when executed by a processor, cause the processor to perform a method. For example, the SmartSwitch includes a PCIe card that is installed in any server to implement distributed services and the SmartSwitch is a top-of-rack switch that also supports distributed services for all the hosts in the rack. https://www.amd.com/system/files/documents/pensando-smartswitches.pdf.

78. The Accused Products collect, at a first endpoint device coupled to a network, parameter values for determination of bandwidth of network connections to further endpoint devices. For example, the SmartSwitch collects parameter values for determination of bandwidth of network connections to further endpoint devices via the servers which collect network telemetry. https://www.amd.com/system/files/documents/pensando-smartswitches.pdf. For example, the Pensando DPUs can determine the transmission bandwidth and other telemetry related to a network flow.

> The Pensando flow-based engine stores each flow in a flow table. Each flow table entry includes counters and timers. For the first time in the history of networking, it is possible to provide precise information on how each flow behaves, including parameters like bytes and packet exchange, delays, jitter, and bandwidth. Previously this information was available only at the aggregate level, for example, at the interface level. This new granularity is what we call "Telemetry," It allows us to pinpoint performance issues at the network and application level.

https://www.amd.com/system/files/documents/pensando-smartswitches.pdf.

79. The Accused Products determine at a tuner server, distinct from the first endpoint device and the further endpoint devices, that a next network connection from the first endpoint device to a second endpoint device matches a geographical area of a past network connection. For example, the SmartSwitch running a Pensando DPU (an example of a tuning server) determines that a next network connection from the endpoint (e.g., a server) to a second endpoint device (e.g., another server) matches a geographical area of past network connection (e.g., flow state table).

The stateful forwarding mechanism is also called *cache-based forwarding*. It relies on a flow cache, a binary data structure capable of "exact" matching packets belonging to a particular flow. The word "exact" implies a binary match easier to implement both in hardware or software, unlike a ternary match, such as LPM. The flow cache contains an entry for each flow, i.e., two entries per session. The flow can be defined with an arbitrary number of fields, thus supporting IPv4 and IPv6 addresses, different encapsulations, policy routing, and firewalling.

Cache entry contains information needed to forward the packet (e.g., layer 2 and/or layer 3 address of the next hop, type of encapsulation, etc.).

The separate initial processing that leads to the forwarding decision may create new cache entries. When a packet is received by the switch, if it does not match any entry in the flow cache (a "miss") it is processed according to the initial processing. Otherwise ("hit"), it is forwarded according to the information in the matching flow cache entry.

This hardware is the *Pensando flow-based engine*, which:

- identifies and tracks all the flows for all the sessions.
- uses this information to forward packets and to apply services
- operates at wire rate
- tracks million of flows simultaneously

https://www.amd.com/system/files/documents/pensando-smartswitches.pdf.

80. The Accused Products initiate the next network connection, from the first endpoint device to the second endpoint device, based on the determination at the tuner server, with a transmission bandwidth based on the parameter values for the past network connection. For example, Pensando DPUs are able to determine the transmission bandwidth and other telemetry related to a network flow.

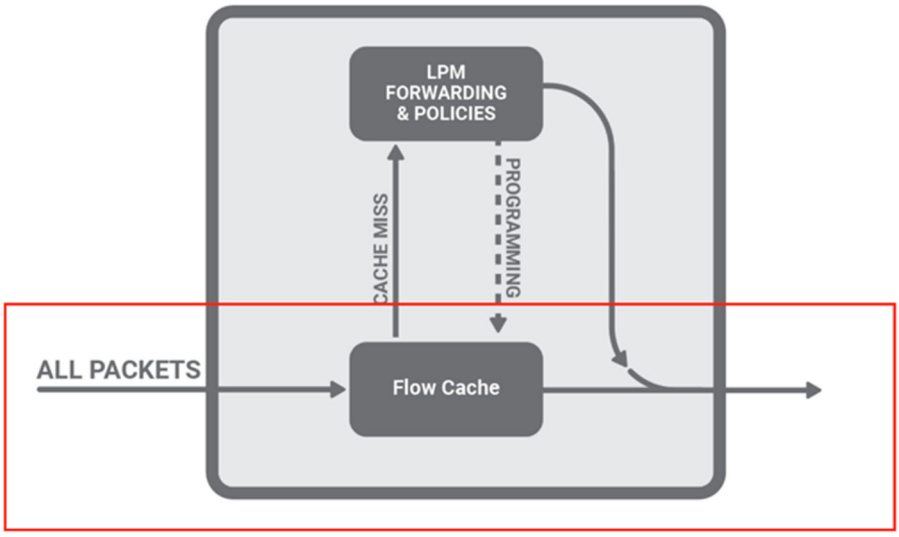

https://www.amd.com/system/files/documents/pensando-smartswitches.pdf.

> The Pensando flow-based engine stores each flow in a flow table. Each flow table entry includes counters and timers. For the first time in the history of networking, it is possible to provide precise information on how each flow behaves, including parameters like bytes and packet exchange, delays, jitter, and bandwidth. Previously this information was available only at the aggregate level, for example, at the interface level. This new granularity is what we call "Telemetry," It allows us to pinpoint performance issues at the network and application level.

https://www.amd.com/system/files/documents/pensando-smartswitches.pdf; *see* https://www.varindia.com/news/pensando-brings-hyperscale-dpu-technology-to-amds-data-center-capabilities ("SmartSwitch technology, as exemplified by the Aruba CX 10000 with AMD Pensando, allows for 'East-West' traffic within the data center itself to be gathered and analyzed, either via intrinsic capabilities of the AMD Pensando Distributed Services Platform, or by providing the data to 3rd party Extended Detection and Response (XDR) vendor products.").

81. On information and belief, Defendants' infringement through its use of its DPU technology and DPU-enabled systems, described above, is exemplary of its infringement with respect to all the Accused Products.

82. Defendants have also indirectly infringed, and continue to indirectly infringe, the '634 Patent under 35 U.S.C. § 271(b) and (c).

83. Defendants knowingly and intentionally actively aided, abetted, and induced others to directly infringe at least claim 8 of the '634 Patent (such as its customers in this District and

throughout the United States), and continue to do so, by, for example, selling and offering access to and encouraging and supporting use of the Accused Products.

84. Defendants contributed to the direct infringement of at least claim 8 of the '634 Patent under 35 U.S.C. § 271(c), and continue to do so, by, for example, supplying, with knowledge of the '634 Patent, a material part of a claimed invention, where the material part is not a staple article of commerce and is incapable of substantial noninfringing use. For example, Defendants have provided, owned, operated, sold, offered to sell, leased, licensed, used, and/or imported, and continue to do so, various hardware and/or software that make up and enable the Accused Products, including as used in third-party (including customer) systems (including as discussed above), are a material part of the claimed invention, are not a staple article of commerce, and are incapable of substantial non-infringing uses.

85. As explained above, Defendants' infringement has been and continues to be willful in view of the facts asserted above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

X.
FIFTH CLAIM
(Infringement of the '767 Patent)

86. Plaintiffs incorporate by reference paragraphs 1-85 herein.

87. On information and belief, Defendants have infringed and continue to infringe the '767 Patent in violation of 35 U.S.C. § 271(a), either literally or through the doctrine of equivalents, by making, using, selling, offering for sale, and/or importing into the United States products and/or methods that practice at least claim 1 of the '767 Patent, including with the Accused Products.

88. Each of the Accused Products meet every limitation of claim 1 of the '767 Patent, which recites:

> 1. A routable packet-switched network supported by an abstraction protocol comprising:
>
> at least one host;
>
> a master controller, the master controller comprising a master processing element, a master controller interface to communicate with the host, and at least one master controller control link interface; and
>
> at least one edge controller, the edge controller comprising an edge processing element, an edge controller control link interface to communicate with the master controller via the master controller control link interface, and at least one storage media device interface to communicate with at least one storage media device,
>
> wherein the master controller and the edge controller communicate via the abstraction protocol, the abstraction protocol comprising a full-duplex protocol supporting full command queuing for the at least one storage media device; and
>
> wherein the edge controller further comprises at least one other master controller further in communication with at least one other edge controller, when the edge processing element associated with the edge controller authors abstract protocol messages intended for the other edge controller.

89. The Accused Products, including AMD Pensando's DPU-enabled systems, make up a routable packet-switched network supported by an abstraction protocol.

> AMD Pensando DPUs are programmed using the industry-standard P4 language, facilitating the implementation of a wide variety of system solutions. It supports software-defined networking and storage protocols, including NVMe virtualization and transport, and is designed to give developers the agility to develop and deploy new features and modifications throughout the product lifecycle.
>
> The Elba DPU form factor and power profile are designed to support multiple system level implementations ranging from a half-height, half-length PCIe card that can fit into the power and cooling profiles of any standard server to network and security appliances and SmartSwitches. The P4-programmable design enables these applications to dynamically re-configure the data processing inside the DPU.

Key Features

Feature	Description
Network Interfaces	56 Gbps PAM4 SerDes supporting Dual 200 GE, Quad-100/50/25/10GE 1GE management port
PCIe Interface	32 Lanes of PCIe Ge n4, configurable as root complex or end-point mode 2x16 / 4x8 / 8x4 with QOS support in multi-host applications
Data Pipeline	P4 data pipeline comprising 144 match processing units (MPUs) @2 GHz Provides high performance capabilities in packet and message processing, at line rate

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf; https://www.amd.com/system/files/documents/pensando-giglio-product-brief.pdf.

90. The Accused Products, including AMD Pensando's DPU-enabled systems, include at least one host.

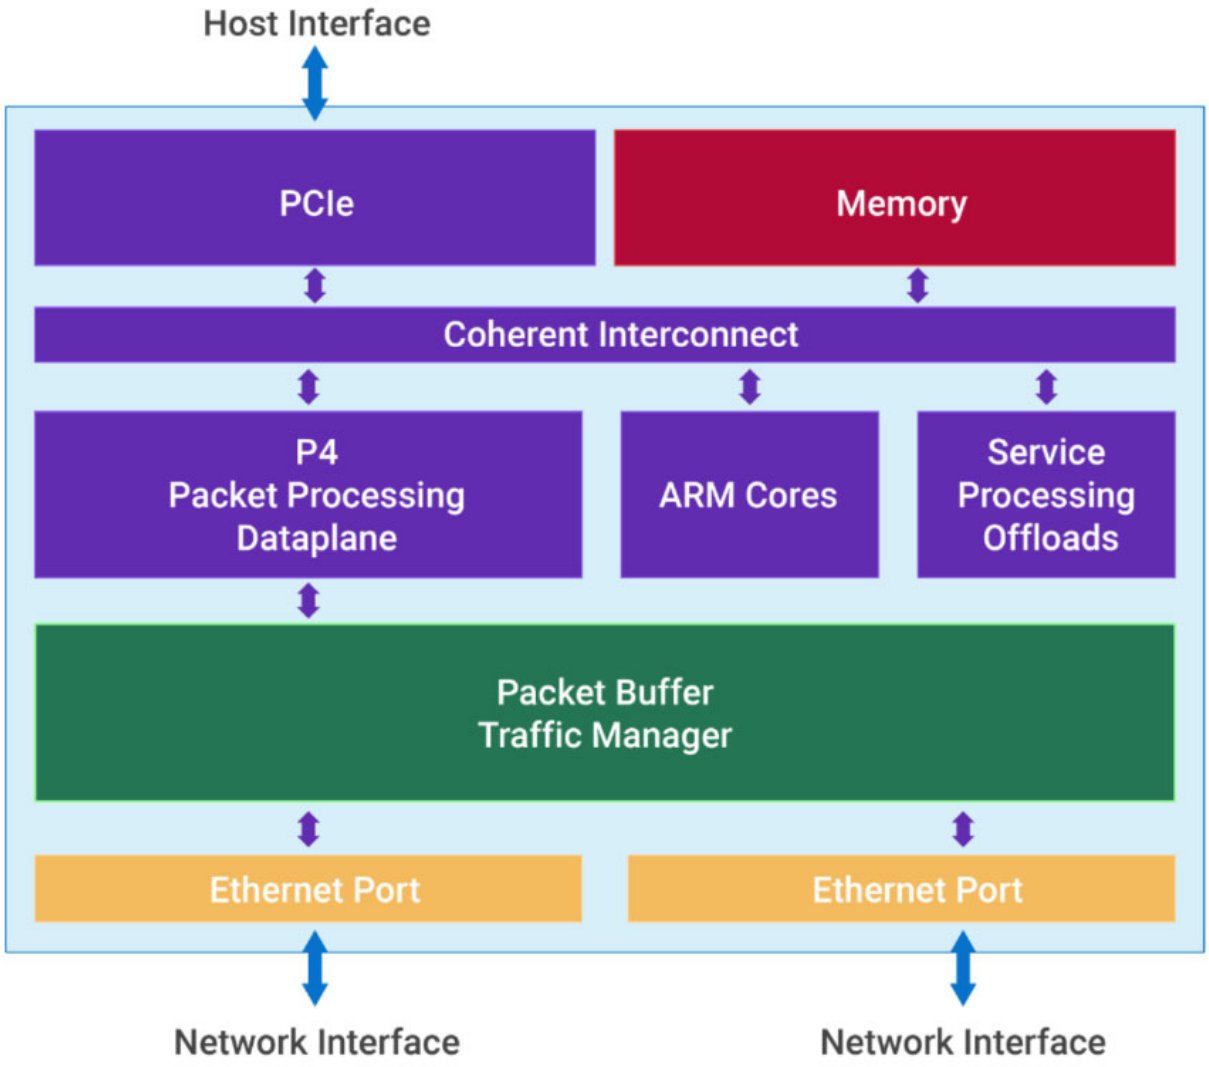

https://infohub.delltechnologies.com/en-US/l/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-poweredge-servers/amd-pensando-dpu.

91. The Accused Products, including AMD Pensando's DPU-enabled systems, include a master controller, the master controller comprising a master processing element, a master controller interface to communicate with the host, and at least one master controller control link interface. For example, the DPU itself can serve as a master controller, and includes a master processing element (e.g., a processor).

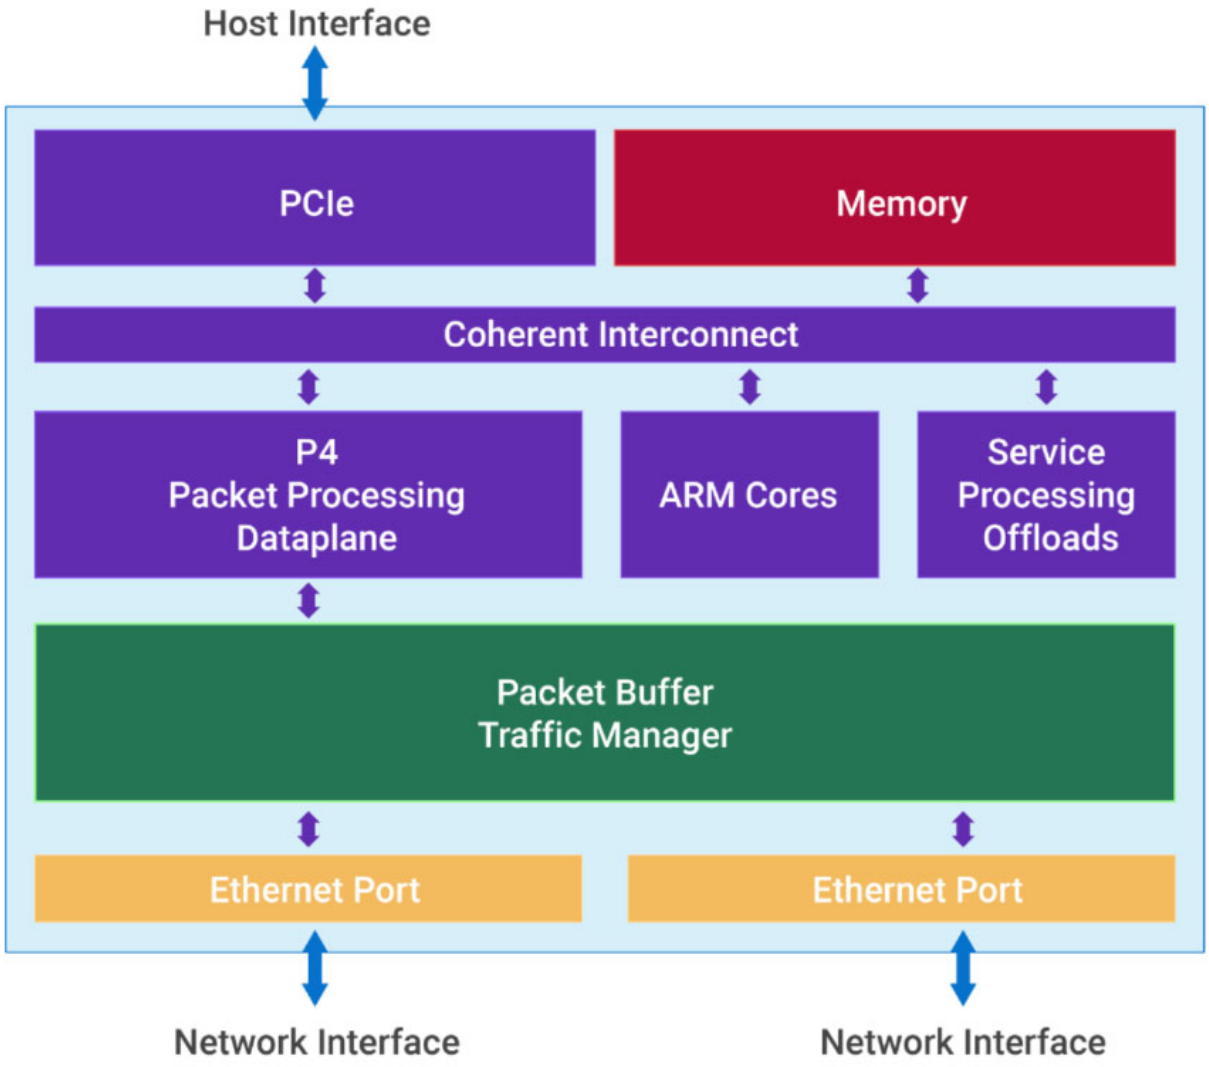

https://infohub.delltechnologies.com/en-US/l/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-poweredge-servers/amd-pensando-dpu.

92. The Accused Products, including AMD Pensando's DPU-enabled systems, include at least one edge controller, the edge controller comprising an edge processing element, an edge controller control link interface to communicate with the master controller via the master controller control link interface, and at least one storage media device interface to communicate with at least one storage media device. For example, the "edge controllers" can be other DPUs, which the master controller can communicate with. In turn, the "edge processing element" can be the processer (e.g., DPU chip) within that edge DPU and the "edge controller control link interface"

is the interface for the network connection that attaches the edge controller to the master controller, with a logical separation for control functions.

Key Applications

Application	Details
Advanced Networking	Full support for SDN. Virtual Private Networks (Network Overlays), L3 ECMP, Load Balancing, NAT, PAT
Cutting-Edge Security Features	Stateful Firewall, Security Groups, Stateless and Reflexive ACLs, VPN Termination (IPsec), TLS/DTLS encryption, TLS Proxy
Enhanced Storage	Full support for SDS. NVMe Virtualization, NVMe-oF with RDMA or TCP Transport, AES-XTS data-at-rate encryption, Compression/Decompression, SHA-3 deduplication, CRC64/32 and checksum acceleration

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf; https://www.amd.com/system/files/documents/pensando-giglio-product-brief.pdf.

93. The Accused Products, including AMD Pensando's DPU-enabled systems, include the master controller and the edge controller communicating via the abstraction protocol, the abstraction protocol comprising a full-duplex protocol supporting full command queuing for the at least one storage media device. *See, e.g.*, https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/solution-brief/csp-case-study.pdf.

> AMD Pensando DPUs are programmed using the industry-standard P4 language, facilitating the implementation of a wide variety of system solutions. It supports software-defined networking and storage protocols, including NVMe virtualization and transport, and is designed to give developers the agility to develop and deploy new features and modifications throughout the product lifecycle.
>
> The Elba DPU form factor and power profile are designed to support multiple system level implementations ranging from a half-height, half-length PCIe card that can fit into the power and cooling profiles of any standard server to network and security appliances and SmartSwitches. The P4-programmable design enables these applications to dynamically re-configure the data processing inside the DPU.

Key Features

Feature	Description
Network Interfaces	56 Gbps PAM4 SerDes supporting Dual 200 GE, Quad-100/50/25/10GE 1GE management port
PCIe Interface	32 Lanes of PCIe Ge n4, configurable as root complex or end-point mode 2x16 / 4x8 / 8x4 with QOS support in multi-host applications
Data Pipeline	P4 data pipeline comprising 144 match processing units (MPUs) @2 GHz Provides high performance capabilities in packet and message processing, at line rate

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf; https://www.amd.com/system/files/documents/pensando-giglio-product-brief.pdf.

Key Applications

Application	Details
Advanced Networking	Full support for SDN. Virtual Private Networks (Network Overlays), L3 ECMP, Load Balancing, NAT, PAT
Cutting-Edge Security Features	Stateful Firewall, Security Groups, Stateless and Reflexive ACLs, VPN Termination (IPsec), TLS/DTLS encryption, TLS Proxy
Enhanced Storage	Full support for SDS. NVMe Virtualization, NVMe-oF with RDMA or TCP Transport, AES-XTS data-at-rate encryption, Compression/Decompression, SHA-3 deduplication, CRC64/32 and checksum acceleration

https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf; https://www.amd.com/system/files/documents/pensando-giglio-product-brief.pdf.

94. The Accused Products, including AMD Pensando's DPU-enabled systems, include the edge controller further comprising at least one other master controller further in communication with at least one other edge controller, when the edge processing element associated with the edge controller authors abstract protocol messages intended for the other edge controller. For example, on information and belief, AMD's DPU edge controllers can contain an element acting as a master controller for communication with other edge controllers. *See, e.g.*, https://www.datanami.com/this-just-in/amd-showcases-continued-enterprise-data-center-

momentum-with-epyc-cpus-and-pensando-dpus ("Additionally, as data center applications continue to grow in scale and complexity, AMD Pensando DPUs are the ideal solution to help offload infrastructure services from the CPU to free up valuable work cycles. VMware vSphere 8-enabled systems, powered by AMD EPYC CPUs and Pensando DPUs, deliver the performance, efficiency and flexibility IT leaders need to run a broad set of business-critical workloads."); https://www.techinsights.com/blog/amds-new-pensando-dpu-cuts-power; https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/solution-brief/csp-case-study.pdf.

95. On information and belief, Defendants' infringement through its use of its DPU technology and DPU-enabled systems, described above, is exemplary of its infringement with respect to all the Accused Products.

96. Defendants have also indirectly infringed, and continue to indirectly infringe, the '767 Patent under 35 U.S.C. § 271(b) and (c).

97. Defendants knowingly and intentionally actively aided, abetted, and induced others to directly infringe at least claim 1 of the '767 Patent (such as its customers in this District and throughout the United States), and continue to do so, by, for example, selling and offering access to and encouraging and supporting use of the Accused Products.

98. Defendants contributed to the direct infringement of at least claim 1 of the '767 Patent under 35 U.S.C. § 271(c), and continue to do so, by, for example, supplying, with knowledge of the '596 Patent, a material part of a claimed invention, where the material part is not a staple article of commerce and is incapable of substantial noninfringing use. For example, Defendants have provided, owned, operated, sold, offered to sell, leased, licensed, used, and/or imported, and continue to do so, various hardware and/or software that make up and enable the Accused Products,

including as used in third-party (including customer) systems (including as discussed above), are a material part of the claimed invention, are not a staple article of commerce, and are incapable of substantial non-infringing uses.

99. As explained above, Defendants' infringement has been and continues to be willful in view of the facts asserted above and its failure to take any action, even after being put on notice, to stop its infringement or inducement of, or contribution to, infringement by others.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), Plaintiffs hereby demand a trial by jury on all issues triable to a jury.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs pray for judgment against Defendants as follows:

A. That Defendants have infringed each of the Asserted Patents, and unless enjoined, will continue to infringe one or more of the applicable Asserted Patents;

B. That Defendants' infringement of one or more of the applicable Asserted Patents has been willful;

C. That Defendants pay Plaintiffs damages adequate to compensate Plaintiffs for Defendants' past infringement of each of the Asserted Patents, and present and future infringement of the applicable Asserted Patents, together with interest and costs under 35 U.S.C. § 284;

D. That Defendants pay prejudgment and post-judgment interest on the damages assessed;

E. That Defendants pay Plaintiffs enhanced damages pursuant to 35 U.S.C. § 284;

F. That Defendants be enjoined from infringing the applicable Asserted Patents, or if its infringement is not enjoined, that Defendants be ordered to pay ongoing royalties to Plaintiffs for any post-judgment infringement of the applicable Asserted Patents;

G. That this is an exceptional case under 35 U.S.C. § 285; and that Defendants pay Plaintiffs' attorneys' fees and costs in this action; and

H. That Plaintiffs be awarded such other and further relief, including equitable relief, as this Court deems just and proper.

Dated: March 29, 2024

Respectfully Submitted,
/s/ *Amy Ruhland*

Michael Matulewicz-Crowley (*pro hac vice* to be filed)
mmatulewicz-crowley@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN
& FELDBERG LLP
400 Madison Avenue, Suite 14D
New York, NY 10017
Telephone: (212) 381-1965
Facsimile: (650) 560-3501

Christine E. Lehman (*pro hac vice* to be filed)
clehman@reichmanjorgensen.com
Ariane S. Mann (*pro hac vice* to be filed)
amann@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN
& FELDBERG LLP
1909 K Street, NW, Suite 800
Washington, DC 20006
Telephone: (202) 894-7310
Facsimile: (650) 560-3501

Amy L. Ruhland (TX Bar No. 24043561)
aruhland@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN
& FELDBERG LLP
901 S. Mopac Expressway, Building 1, Suite 300
Austin, TX 78746
Telephone: (650) 623-1401
Facsimile: (650) 560-3501

Attorneys for Plaintiffs
Concurrent Ventures, LLC and
XtreamEdge, Inc.