

Exhibit D: Claim Chart for U.S. Patent 10,985,943

U.S. Patent No. 10,985,943

This claim chart is based on publicly available literature, the source code¹ Defendants have made available, and the documents currently produced by Defendants (AMD-CV-0000001–AMD-CV-00005452). Plaintiffs Concurrent Ventures, LLC and XstreamEdge, Inc. (“Plaintiffs” or “CV”) reserve the right to supplement and/or amend their positions based on discovery from Defendants Advanced Micro Devices, Inc. and Pensando Systems, Inc. (“AMD Pensando” or “Defendants”), including technical documentation such as schematics and datasheets, and source code. Plaintiffs reserve the right to amend these contentions for good cause with leave of the Court.

Defendants infringed and continue to infringe U.S. Patent No. 10,985,943 (the “’943 patent”) by having made, used, sold or offered for sale, and/or imported into the United States their Distributed Services Cards (“DSCs”) and Data Processing Units (“DPUs”) including at least Elba, Capri, Giglio, and Salina (including AMD DSC2-25G, AMD DSC2-100G, and AMD DSC2-200), as well as their DPU and DSC technology when implemented in other products, services, and offerings, for example the SmartSwitch or AMD Pensando Pollara 400² (collectively, the “Accused Products,” “Accused DPUs,” or “Accused DSCs”). For an example, Defendants directly infringe when they themselves implement, use, or test the Accused Products in the claimed manner discussed herein.

Additionally, Defendants also indirectly infringe the ’943 patent by having induced and/or contributed to third parties’ infringement, including the infringement of Defendants’ customers. In particular, Defendants, through their advertising, marketing, communications with customers and clients, troubleshooting, guidance, testing, setting of default settings, and helping, actively and knowingly aids and abets their customers’ and clients’ infringement, including their use of the Accused Products. Defendants also sold or offered to sell into the United States, or imported into the United States, their Accused Products, which are material components of the infringing implementation and not staple articles of commerce capable of substantial noninfringing use, as shown, for example, at least by Defendants’ own descriptions, guidance, and instructions to implement the Accused Products. Defendants have known of the patents, their infringement, and that the Accused Products are especially made for use in a manner that would infringe the patent when they sold, offered to sell, and/or imported the Accused Products. Further details are provided herein.

Theories exemplifying how the Accused Products infringe the asserted claims of the ’943 patent are detailed herein. The products included are representative; on information and belief, all Accused Products function in substantially the same manner as one or more

¹ Due to the parties’ ongoing negotiations regarding the protective order and the time constraints inherent in the draft procedures for printing source code in the protective order, these contentions incorporate Plaintiffs’ review of the source code but do not explicitly refer to the source code.

² <https://www.tomshardware.com/networking/amd-unveils-industrys-first-ultra-ethernet-ready-network-card-for-ai-and-hpc>, CV-AMD0016365; CV-AMD0017532.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

of the examples provided, and include generally the same components. For example, for purposes of infringement of the '943 patent, Defendants' Elba, Capri, Giglio, and Salina versions operate, in relevant part, in the same manner, and thus infringe in substantially the same way. *See* AMD-CV-0002195–AMD-CV-0002196; AMD-CV-0004986–AMD-CV-0005004 at AMD-CV-0004988; AMD-CV-0003945–AMD-CV-0004184 at AMD-CV-0003962; AMD-CV-0002195; AMD-CV-0004986; Defendants' Response to Interrogatory Nos. 1, 2. The following documents provide additional support for Plaintiffs' infringement theories documented herein: AMD-CV-0001962; AMD-CV-0000001; AMD-CV-0000608; AMD-CV-0000643; AMD-CV-0000935; AMD-CV-0000988; AMD-CV-0000994; AMD-CV-0001126; AMD-CV-0001155; AMD-CV-0002797; AMD-CV-0002200; AMD-CV-0004235; AMD-CV-0003945–AMD-CV-0004184; Defendants' Response to Interrogatory Nos. 1, 2. This is further shown by the evidence for the various versions cited herein.

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
[1pre] A programmable device for data flow processing in a user-configurable server with swappable pods or cards, comprising:	To the extent the preamble is limiting, Defendants make, use, test, sell, and offer to sell the claimed programmable device for data flow processing in a user-configurable server with swappable pods or cards, including by making, using, testing, selling, and offering to sell the Accused Products; implementing and/or testing those Accused Products in their own platforms; and providing them to customers. See, e.g., AMD-CV-0003945–AMD-CV-0004184. AMD Pensando Accused Products (AMD Distributed Service Card (DSC) or AMD Pensando Distributed Modules (DSMs)) are Data Processing Units (DPUs) and enabled system and programmable devices for data flow processing in a user-configurable server with swappable pods or cards. Further, AMD Pensando DSCs are deployed in Server, Smart Switch, and Appliance platforms for data flow processing. The AMD Pensando Data Processing Unit (DPU) features user-configurable P4 Pipelines for dynamic data processing.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Cloud Solution Form Factor Form Factors  DPU in Server  DPU in Appliance  DPU in SmartSwitch https://nowlab.cse.ohio-state.edu/static/media/workshops/presentations/exacomm23/exacomm2023_final-Kimon.pdf, CV-AMD0016028.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD SECURING THE DATA CENTER COMPREHENSIVE SECURITY FROM CORE TO EDGE Distributed Security Services Extend the secure perimeter Edge and Beyond Smart Switch & Appliance Secure east-west traffic and legacy endpoints Infrastructure Advanced SmartNIC Secure every endpoint Confidential Compute Secure the application environment Root of Trust Secure the CPU at boot Compute Node Advanced Telemetry Monitor and detect threats in real time AMD EPYC, PENSANDO, ALVEO, VERSAL elastic, FORINET, Guardicore, splunk https://www.servethehome.com/amd-data-center-solutions-at-amd-financial-analyst-day-2022/amd-fad-2022-amd-securing-the-data-center/, CV-AMD0015776.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	DSC2-25G The AMD Pensando Architecture Delivers  HARDWARE SPECIFICATIONS Form Factor: Half height, Half-length PCIe PCIe Interface: 8-lane PCIe Gen 4.0 Network Ports: 2 ports SFP56 Port Configuration: 2x10G, 2x25G Processor: Elba 7nm NETWORKING, SECURITY and STORAGE SERVICES Advanced Observability: Flow-based packet telemetry, stateful connection statistics, latency metrics, drop statistics, threshold alerting, ERSPAN (bi-directional), NetFlow/IPFIX Advanced Networking: Virtual Private Networks (network overlays), L3 ECMP, Load balancing, NAT, PAT Advanced Security: Stateful firewall with connection tracking, security groups, stateless and reflexive ACLs, VPN termination (IPsec), TLS/DTLS encryption, TLS Proxy Enhanced Storage: NVMe virtualization, NVMe-oF with RDMA or TCP transport, AES-XTS data-at-rest encryption, compression, SHA-3 deduplication, CRC64/32 acceleration <i>AMD Pensando DSC2 25G</i> https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015842–CV-AMD0015843.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	There is also a DSC2-100G. This is the dual QSFP28 card for 2x 100GbE. DSC2-100G The AMD Pensando Architecture Delivers  HARDWARE SPECIFICATIONS Form Factor: Full height, Half-length PCIe PCIe Interface: 16-lane PCIe Gen 4.0 Network Ports: 2 ports QSFP28 Port Configuration: 2x100G Processor: Elba 7nm PERFORMANCE AND SCALE Delivers wire-speed services on each of its QSFP28 ports, including chained services such as L4 stateful firewall + IPsec encryption + load balancing NETWORKING, SECURITY and STORAGE SERVICES Advanced Observability: Flow-based packet telemetry, stateful connection statistics, latency metrics, drop statistics, threshold alerting, ERSPAN (bi-directional), NetFlow/IPFIX Advanced Networking: Virtual Private Networks (network overlays), L3 ECMP, Load balancing, NAT, PAT Advanced Security: Stateful firewall with connection tracking, security groups, stateless and reflexive ACLs, VPN termination (IPsec), TLS/DTLS encryption, TLS Proxy Enhanced Storage: NVMe virtualization, NVMe-oF with RDMA or TCP transport, AES-XTS data-at-rest encryption, compression, SHA-3 deduplication, CRC64/32 acceleration AMD Pensando DSC2 100G https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015843.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	SUPERCHARGING OUR PORTFOLIO Cloud Smart Appliance - Expand Existing Cloud Footprint - Rinse & Repeat with Other Clouds Cloud DPU's - Double Down On Existing Deployments - Expanding into MDC's Enterprise DPU's - AMD the Project Monterey Platform of Choice - AMD EPYC Solutions (Future) Smart Switch - Win More Logos Logos: Dell Technologies, Hewlett Packard Enterprise, IBM, Microsoft, ORACLE, vmware, AMD together we advance. 11 AMD VMware Explore Press Briefing Under embargo until 8am ET on August 30, 2022 AMD Pensando Portfolio 2022 08 30 https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015843. <i>The Pensando distributed services architecture is designed to provide programmable security, network, storage, and visibility services to a data center landscape growing in scale and complexity. This domain-specific architecture combines P4 packet processing pipelines with embedded ARM CPUs and hardware offload engines in a tightly coupled, cache coherent interconnect. Millions of flows are tracked by hardware and distributed to a combination of P4 programs, ARM processes, custom hardware engines, and PCIe based hosts and devices. There are currently 16- and 7-nm ASIC implementations of this distributed services architecture, future generations are in design.</i> Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016133.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	P4 Pipelines Pensando ASICs include multiple P4³ 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 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 Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016134.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Pensando Announces P4-programmable Platform and Joins P4 Community FEB 24, 2020 Mario Baldi ABOUT THE AUTHOR Posted by Mario Baldi on February 24, 2020 Pensando Systems recently introduced a Distributed Services Card (DSC) based on a custom, fully-programmable processor optimized to execute P4 programs. The DSC is the centerpiece of a Distributed Services Platform designed to deliver high-performance, scalable networking, storage and security services. This unique approach removes the appliance bottleneck, offers unmatched visibility on traffic, and provides significant benefits for various use cases in private, public or hybrid clouds, as well as enterprise data center networks where east-west security concerns are increasingly becoming an issue. Easily installed in standard servers, the DSC, shown in the figure below, includes specialized hardware blocks to accelerate computation intensive functions, such as compression and encryption, and relies on P4 to provide dataplane programmability at wire rate. 

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Pensando Announces P4-programmable Platform and Joins P4 Community - Open Networking Foundation by Mario Baldi (2020), CV-AMD0016130; https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015842. AMD Pensando DSCs are deployed in server, for example, as shown below, server platforms provided by Dell, HPE, Microsoft, Oracle, include AMD Pensando DSCs to provide data flow processing. Dell servers showcase world-record performance, scale dynamically to run highly demanding applications and workloads, and in-built features such as cyber resiliency, automation, sustainability, and global services. Dell Technologies in collaboration with VMware on their Project Monterey initiative, and by partnering with Nvidia and AMD has incorporated these technologies onto the PowerEdge server portfolio to deliver an end-to-end integrated solution to enterprises. • AMD/Pensando – There are two offerings from AMD/Pensando: ▪ AMD DSC2-25G DPU with 2 ports of 10/25G based on 2nd-generation P4-programmable SoC with 16 ARM cores. ▪ AMD DSC2-100G DPU with 2 ports of 100G based on 2nd-generation P4-programmable SoC with 16 ARM cores. https://infohub.delltechnologies.com/en-us/t/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-powerededge-servers/, CV-AMD0016336–CV-AMD0016337. • AMD Is the Clear Platform of Choice for vSphere 8: Of the three DPU vendors chosen by VMware for the vSphere Distributed Services Engine (or vSphere on DPUs), AMD is the <i>only</i> vendor with supported solutions from both vDSE launch server partners (Dell and HPE). • AMD Offers More Flexibility: AMD is the only vendor at launch with 10G/25G and 100G solutions on the market. • AMD Is the Only Proven DPU Platform: AMD is the only vendor with large scale of DPU units deployed in production with hyperscalers, public clouds, and enterprise customers. https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/white-papers/vdse-with-amd-white-paper.pdf, CV-AMD0015898.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Support from AMD Pensando Partners  AMD Pensando Distributed Services Cards are available as factory-ordered options on select HPE ProLiant, HPE Apollo, and HPE Edgeline platforms in addition to availability through HPE GreenLake, HPE's market-leading as-a-service platform. Visit HP   AMD Pensando Distributed Services Cards are available as factory-ordered options on select Dell products. Visit Dell  https://www.amd.com/en/support/accelerators/pensando.html, CV-AMD0015777. PENSANDO DISTRIBUTED SERVICES PLATFORM FOR HPE SERVERS https://www.hpe.com/psnow/doc/a50001239enw, CV-AMD0016288.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	“Performance, efficiency and security are integral components of competitive offerings within cloud computing and hyperconverged infrastructure,” said Forrest Norrod, senior vice president and general manager, Data Center Solutions Group, AMD. “VMware vSphere 8, running on AMD Pensando DPUs, is a key step in bringing the industry closer to composable hardware systems and truly pervasive heterogeneous computing. AMD is uniquely positioned in the industry to deliver performance, efficiency and security with a distributed services platform that is currently deployed with major customers including Goldman Sachs, Microsoft Azure, NetAPP, IBM Cloud and Oracle Cloud.” Industry-Leading AMD Pensando DPUs Enable Accelerated Data Centers with VMware vSphere 8: Advanced Micro Devices, Inc. (AMD), CV-AMD0016004. As IBM began shifting its focus to scalable cloud-native architecture in 2019, it sought out infrastructure that could scale to the cloud’s needs by providing high levels of performance, security and architecture resiliency. The vendor found its answer in AMD Pensando’s data processing unit (DPU). https://www.sdxcentral.com/articles/interview/why-ibm-picked-amd-pensando-dpus-for-scalable-cloud-native-architecture/2023/07/, CV-AMD0015874. Accelerating Microsoft Azure with AMD DPUs  Kevin Pacella Staff 3 0 8,569 05-01-2023 10:40 AM I’m pleased to share with you that AMD Pensando DPUs are accelerating Microsoft Azure. With the recent announcement of <i>Accelerated Connections</i> available now for Public Preview, Azure customers can achieve orders of magnitude higher networking performance. This new service offering is an architectural shift that decouples the DPU adapter and SDN stack from the server. https://community.amd.com/t5/corporate/accelerating-microsoft-azure-with-amd-dpus/ba-p/604170, CV-AMD0015793.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Microsoft Azure “Pensando System’s full stack solution, complete with a P4 programmable processor and hyper scale software has enabled Microsoft Azure to rapidly augment its networking and security services improving performance, latency, and scale,” said Girish Bablani, corporate vice president, Microsoft Azure. “We have seen a 40x improvement in overall cloud based connection related performance. Pensando delivered this in less than 12 months. This partnership has enabled a transformational “limitless networking” shift that will ensure our customers have access to the highest performing products and services they expect from Microsoft Azure.” https://ir.amd.com/news-events/press-releases/detail/1057/amd-expands-data-center-solutions-capabilities-with, CV-AMD0016153. Oracle Cloud “By partnering with Pensando, Oracle Cloud Infrastructure (OCI) has developed an industry leading Software Defined Networking stack, delivering infrastructure features to Oracle Cloud customers, that have the best price to performance characteristics in the cloud industry,” said Clay Magouyrk, executive vice president, Oracle Cloud Infrastructure. “Pensando’s P4 programmable platform enables OCI to continuously innovate and deliver secure and inexpensive network services with high velocity. The ability to rapidly add new features using a programmable hardware platform will help OCI leapfrog its cloud competitors. Pensando’s engineering team, consisting of many industry experts, has collaborated very successfully with OCI engineering, and have operated at times as OCI’s extension team. OCI views Pensando as a key strategic partner in achieving our goal of being the best infrastructure platform in the cloud industry.” https://ir.amd.com/news-events/press-releases/detail/1057/amd-expands-data-center-solutions-capabilities-with, CV-AMD0016153. Furthermore, AMD Pensando DSCs are deployed in OEM vendor Smart Appliance and SmartSwitch platforms, including the HPE Aruba CX 10000 Switch Series.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	SmartSwitches are fundamental building blocks of a Distributed Services Platform in a top-of-rack form factor. This new category of product can be easily integrated into any new or existing enterprise and cloud environment, regardless of whether it is public, private or hybrid. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016041. Elba DPUs are also integrated into OEM vendor Smart Appliance and Smart Switch platforms, such as the HPE Aruba CX 10000 Switch Series. See the below product data sheets for product overviews, functionality and ordering specifics. • HPE Aruba Networking CX 10000 Switch Series ◦ Data Sheet ◦ QuickSpecs https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf, CV-AMD0015774.  The switch combines the best-of-breed network operating system, Aruba AOS-CX, for data center, campus, and edge in addition to the fully programmable AMD Pensando DPU. This allows the CX 10000 to deliver stateful software-defined services inline, at scale, with wire-rate performance and orders of magnitude scale and performance improvements over traditional L2/3 switches at a fraction of their TCO. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/data-sheets/aruba-cx-10000-series-datasheet.pdf, CV-AMD0015986.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
[1a] a programmable logic device (PLD) having a first region and a second region;	The Accused Products comprise a programmable logic device (PLD) having a first region and a second region. For example, the AMD Pensando DSC contains a programmable logic device (PLD) having a first region and a second region. The AMD Pensando DSC includes hardware or software modules that are programmable using P4 pipelines to accelerate and customize cloud infrastructure services. Pensando Distributed Services Card • High-speed networking connectivity - Pensando has two 200Gbps MACs which can allow for up to 2x 200GbE networking. • High-speed packet processing with specific acceleration and <u>often programmable logic</u> - <u>The Elba DPU is utilizing a P4 programmable path to align with the popular language in the networking space.</u> • A CPU core complex - Here we a complex again of sixteen Arm A72 core complexes • Memory controllers - Here we get dual-channel DDR4/ DDR5 memory support with 8-64GB. The previous generation utilized HBM, but Pensando is using DDR now to save costs and make the design more flexible since one can use slotted memory as well. • Accelerators - There are additional accelerators cryptography, compression, and data movement, among others. This is beyond the P4 network complex • PCIe Gen4 lanes - We get a up to 32x PCIe Gen4 lanes and 8 ports (for example for 8x x4 NVMe SSDs) • Security and management features - The control plane has featured here such as a hardware root of trust. These chips also have a 1GbE NCSI interface for out-of-band management. • Runs its own OS separate from a host system - Pensando says this runs Linux with DPDK support. VMware is also prioritizing this along with the NVIDIA BlueField-2 as VMware Project Monterey. Key Focus: Networking <i>Pensando DSC DPU Example Q2 2021</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147.  The block diagram shows the Elba DPU architecture with components like PCIe MAC, P4 P4P, P4 TxDMA, P4 RxDMA, P4 Ingress, P4 Egress, Packet Buffer, and various controllers. Below it is a photograph of the physical hardware card. Further, the PLD included in the accused product has a first region and the second region which include multiple hardware and/or software modules. For example, the Coherent interconnect, P4 Ingress, P4 Egress and packet buffer are some of the modules in the first region. Similarly, P4 TxDMA, P4RxDMA, ARM cores are few examples of the modules within the second region.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Capri (16 nm) Block Diagram The diagram illustrates the Capri (16 nm) block architecture. At the top, a PCIe MAC and VF BAR maps are connected to an NX Coherent Interconnect + 1MB HBM Cache. This interconnect is linked to HBM (High Bandwidth Memory) and a P4 Tx/DMA block. The P4 Tx/DMA block contains a Scheduler, P4 TxDMA, and P4 Rx/DMA. Below this, there are P4 Egress and P4 Ingress blocks, each containing a Packet Buffer Traffic Manager. The P4 Egress block is connected to Ethernet Port 0, and the P4 Ingress block is connected to Ethernet Port 1. The central part of the diagram shows a large green block labeled 'Packet Buffer Traffic Manager'. To the left of this block are various control and management blocks: Register Interface, Interrupt Controller, timers, doorbells, Misc IO, crypto, PKA, compress/decompress, semaphores, and atomic writes. To the right are ARM A72 cores and L2 Cache. The bottom of the diagram shows connections to Ethernet Port 0, NCSI, and Ethernet Port 1. The entire system is labeled 'Capri'. <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 FIGURE 1. Asic block diagram. Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016134. Pensando Stateful Software Services Embedded natively in the Aruba CX 10000 Switch Series, the AOS-CX operating system is an industry unique, inline stateful software service, enabled through Pensando’s <u>programmable DPU</u>, delivered at scale with wire-rate performance. The services include stateful firewalling and secure segmentation, DDoS protection, and deep flow-based session-level telemetry, enabled natively in the fabric switching infrastructure without requiring host-based agents or dedicated appliances. https://www.amd.com/system/files/documents/aruba-cx-10000-series-datasheet.pdf, AMD-CV-0004579.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Aruba CX 10000 with Pensando Help protect the unprotected with a <u>fully-programmable Aruba switching platform optimized to execute the Pensando Distributed Services Data Processing Unit (DPU) and Software Stack</u> Programmability of Pensando DSC  Figure 2: Programmability of the P4-programmable pipelines in the Pensando Distributed Services Card. Source: Principled Technologies, based on "Solution Brief: Distributed Services for Cloud Providers."⁵ Cloud providers can program the P4-programmable data plane in the Distributed Services Card. Figure 2 shows the P4-programmable data pipeline, which network engineers can leverage to customize each layer of their infrastructure stack. https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf, CV-AMD0015867.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 <i>Figure 14. Programmable pipeline</i>
[1b] 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 Accused Products include 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 Accused Products includes a packet buffer, and P4 Ingress and Egress pipelines which work together as a traffic manager and route the packets to different modules. Further, the packet buffer has multiple ports which are connected to P4 pipeline modules including P4 TxDMA, and P4 RxDMA. Also, the first region has NX Coherent Interconnect which enables the communication between the first region and the second region. Each of these, separately and together, constitute the claimed “hardware-based or firmware-based router” and “port.”

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Capri (16 nm) Block Diagram The diagram illustrates the Capri (16 nm) block architecture. At the top, a PCIe MAC and VF BAR maps are connected to the system. A large HBM (High Bandwidth Memory) block is shown on the right. The central NX Coherent Interconnect + 1MB HBM Cache connects to various components. On the left, a Register Interface, Interrupt Controller, timers, doorbells, and Misc IO are shown. A crypto block containing PKA, compress/decompress, semaphores, and atomic writes is also present. The P4 Tx/DMA and P4 Rx/DMA blocks are connected to the interconnect. Below these are P4 Egress and P4 Ingress blocks, each containing a scheduler, sram, HBM cache, and TCAM. The central Packet Buffer Traffic Manager is connected to the P4 Egress and P4 Ingress blocks. The system also includes L2 Cache and four ARM A72 processors. At the bottom, Ethernet Port 0 and Ethernet Port 1 are shown, connected via NCSI and timestamp blocks. The diagram is labeled 'Hot Chips 32 Pensando Capri Block Diagram'. <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible, CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Life of a Packet Packets received through the network enter on one of the Ethernet MAC interfaces and are placed in the packet buffer according to L2 or L3 class of service (COS). The packet buffer provides pause absorption buffering and COS aware arbitration for the P4 ingress pipeline. Consistent with the P4 model, packets pass through the P4 ingress pipeline for flow classification, firewall, tunnel endpoint processing, and other ingress services; return to the packet buffer in case of 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. mmi202102.issue.pdf (gibsonnet.net), CV-AMD0016199–CV-AMD0016200.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 The diagram illustrates a network processing system architecture. At the top, a 'Host Interface' connects to a 'PCIe' block. The 'PCIe' block is connected to a 'Coherent Interconnect' block, which in turn connects to a 'Memory' block. Below the 'Coherent Interconnect' are three blocks: 'P4 Packet Processing Dataplane', 'ARM Cores', and 'Service Processing Offloads'. These three blocks are connected to a 'Packet Buffer Traffic Manager' block. Below the 'Packet Buffer Traffic Manager' are two 'Ethernet Port' blocks. Finally, each 'Ethernet Port' block connects to a 'Network Interface' at the bottom. Bidirectional arrows indicate communication between adjacent components. https://infohub.delltechnologies.com/en-us/1/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-powerededge-servers/amd-pensando-dpu/, CV-AMD0016334.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Figure 13. AMD Pensando Architecture A <i>pipeline</i> is programmed in P4 and runs in the elements that connect to the packet buffer, as depicted in Figure 14. A pipeline can include various stages that may provide one or more functions. Each stage can run in parallel on different MPUs, and as packets move through the pipeline, they are not waiting on clock cycles that may be servicing a different stage. In addition, all the MPUs have access to the same memory system, so that a packet passing from one stage to the next is not being copied to a new memory location. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015807.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	As another example, as shown below, the packet buffer is capable of routing packets to P4 pipelines such as P4 TxDMA and P4 RxDMA. Figure: P4 Permit Action X28794-110623 https://docs.amd.com/r/en-US/ug1669-amd-pensando-ssdk-hello-world-user-guide/Redirect-P4-Action_CV-AMD0016083.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
[1c] the second region comprising one or more sandboxes with user-definable programmable electronic circuits of the PLD,	The Accused Products include the second region comprising one or more sandboxes with user-definable programmable electronic circuits of the PLD. For example, the second region comprises one or more sandboxes with user-definable programmable electronic circuits of the PLD. For example, the second region contains the DMA P4 pipelines (P4 TxDMA, and P4RxDMA) which function as the claimed sandboxes being that users can program them to function as needed. In addition, the accused product also includes ARM cores (second region) that users can program, and also are the claimed sandboxes. Capri (16 nm) Block Diagram  The diagram illustrates the Capri (16 nm) architecture. At the top, it shows x16 Gen3/4 PCIe MAC and VF BAR maps connected to a central Capri block. To the right is HBM (High Bandwidth Memory). Below these is the NX Coherent Interconnect + 1MB HBM Cache. The Capri block contains several P4 pipelines: P4 TxDMA (with Scheduler, sram, HBM Cache, and DMA), P4 RxDMA (with DMA, sram, HBM Cache, and P4V), P4 Egress (with sram, cache, TCAM, and Parser), and P4 Ingress (with sram, cache, TCAM, and Parser). A central Packet Buffer Traffic Manager connects these pipelines. Other components include a Register Interface, Interrupt Controller, timers, doorbells, Misc IO, crypto, PKA, compress/decompress, and semaphores/atomic writes. At the bottom, it shows Ethernet Port 0 and Ethernet Port 1 with timestamp and MAC address inputs, and an NCSI interface. https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc, CV-AMD0015842.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Programmability of Pensando DSC  Figure 2: Programmability of the P4-programmable pipelines in the Pensando Distributed Services Card. Source: Principled Technologies, based on "Solution Brief: Distributed Services for Cloud Providers."⁵ Cloud providers can program the P4-programmable data plane in the Distributed Services Card. Figure 2 shows the P4-programmable data pipeline, which network engineers can leverage to customize each layer of their infrastructure stack. https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf, CV-AMD0015867.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	<i>Figure 14. Programmable pipeline</i> In contrast, the AMD Pensando DPU architecture is purpose-built to accelerate services in silicon. The programmable silicon consists of <i>match processor units</i> (MPUs) capable of executing a variety of services. Using a toolchain that is centered around the P4 language, the AMD Pensando DPU can be programmed to accelerate a variety of services and protocols. Due to its programmable nature, it can also be programmed to handle new protocols or customized frame types that include metadata for special handling. The general-purpose Arm cores in the AMD Pensando architecture are used primarily for control plane and exception packets. Figure 13 shows the architecture of the AMD Pensando DPU. The key central components of the DPU are the MPUs; this is where the software-in-silicon is executed and provides accelerated (fast path) services. Note that the system memory is connected to both the general-purpose cores and the domain-specific MPUs. This allows features that are typically handled by general-purpose cores, like connection tracking, to be implemented on the MPUs and executed directly in the pipeline. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015808.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AGILE PLATFORM FOR CLOUD PROVIDERS The DSC-100 is the ideal software-defined platform to bring high-performance and efficiency to the cloud infrastructure. Developers can create customized data plane services that are targeted to 100G throughput, microsecond-level latencies, and that scale to millions of flows. High-level programming languages (P4, C) enable rapid development and deployment of new features and services. The programmability and agility of the DSC platform allows customized functionality to be developed by cloud providers, speeding innovation and competitive differentiation. • Incorporates both data plane and control plane, eliminating host agents • Customizable data plane - customers can develop their own software on the platform INDUSTRY-LEADING HARDWARE PLATFORM Pensando's P4 Programmable Processor powers the DSC, enabling wire-speed performance – even with chained services – as well as enhanced security through isolated enforcement of policy. The data plane and control planes are fully software-defined and supported with optimized hardware accelerators. Up to 8 GBytes of on-board DRAM provides the performance and flow-table capacity for true cloud-scale deployments. https://cmssc28productionstorage.blob.core.windows.net/article-assets/5eeeb987d386d2c61a5e67b0-2.pdf, CM-AMD0016072.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement  Pensando Distributed Services Card • High-speed networking connectivity - Pensando has two 200Gbps MACs which can allow for up to 2x 200GbE networking. • High-speed packet processing with specific acceleration and often programmable logic - The Elba DPU is utilizing a P4 programmable path to align with the popular language in the networking space. • A CPU core complex - Here we a complex again of sixteen Arm A72 core complexes • Memory controllers - Here we get dual-channel DDR4/ DDR5 memory support with 8-64GB. The previous generation utilized HBM, but Pensando is using DDR now to save costs and make the design more flexible since one can use slotted memory as well. • Accelerators - There are additional accelerators cryptography, compression, and data movement, among others. This is beyond the P4 network complex • PCIe Gen4 lanes - We get a up to 32x PCIe Gen4 lanes and 8 ports (for example for 8x x4 NVMe SSDs) • Security and management features - The control plane has featured here such as a hardware root of trust. These chips also have a 1GbE NCSI interface for out-of-band management. • Runs its own OS separate from a host system - Pensando says this runs Linux with DPDK support. VMware is also prioritizing this along with the NVIDIA BlueField-2 as VMware Project Monterey. Key Focus: Networking https://www.servethehome.com/dpu-vs-smartnic-sth-nic-continuum-framework-for-discussing-nic-types/pensando-dsc-dpu-example-q2-2021/, CV-AMD0016147.
[1d] 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.	The Accused Products include the port 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. For example, the port included in the AMD Pensando DSC's first region has a bridge that is lockable to prevent user access, and unlockable to enable user access, to each of the programmable P4 engines.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
of the sandboxes and the user-definable programmable electronic circuits therein.	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900. Based upon information and belief, the AMD Pensando DPU does not allow user to directly access the P4 engines or the P4 pipelines. To program the P4 engines, the user must log in with SSDK compiler. Accordingly, the port thus includes a bridge that is lockable to prevent user access, and unlockable to enable user access, to each of the sandboxes and the user-definable programmable electronic circuits therein.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Copying an Image to the DSC After building the DSC firmware image you must copy it to the DSC. You can use any of the following in-band or out-of-band management interfaces: • The internal management NIC (int_mnic) interface • The memtun interface • A serial cable connected to the DCS console • The external RJ45 out-of-band (OOB) interface Using the Internal Management NIC Interface To copy the DSC firmware image using the internal management NIC (int_mnic) interface: 1. Boot up the host with the DSC plugged in. 2. <u>Login to the host.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0015832.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando you can log into the DPU, but really the way you are meant to program the part is through the P4 engines.  AMD Pensando DSC2 100 100G 2P QSFP56 DPU Login Top https://www.servethehome.com/hands-on-with-an-amd-pensando-elba-dpu-in-a-secret-lab-arm-nvidia-dell-vmware-esxi-upt/, CV-AMD0016077.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando has various hardware security features to prevent unauthorized access to the P4 engines. For example, AMD Pensando has a hardware based firewall at the PCIe interface which can restrict the user access to the P4 engines. Accordingly, it restricts any user to access the P4 engines as well as ARM cores. Pensando's Security Centric Solution Hardware-Enforced Security The highest level of security assurance is best achieved with isolated hardware that is architected to be secure from all expected threats, both internal and external. Pensando has developed the Capri Programmable Services Processor (PSP), based on a domain-specific architecture that integrates a wide range of security features described in this document. This device is deployed on Pensando's Distributed Services Cards (DSCs) that are installed in data center servers to bring secure services directly to the server edge. HardGap™ Firewall from Host In order to provide isolation from server-based vulnerabilities including deliberate attack or a rogue process writing through memory, reliable protection of host PCIe access to the DSC memory regions is required. Pensando has developed a unique hardware-based firewall at the PCIe-memory interface inside the Programmable Services Processor. Unlike traditional PCIe-attached devices, Pensando Programmable Services Processors provide no intrinsic device access to the host complex. In other words, the host has no de-facto privilege to control the device. The PCIe device tree that the host system sees is established entirely by DSC trusted software, and enforced by hardware PCIe transaction filters. The DSC software controls the PCIe devices presented by the DSC to the host, as well as the individual resource BARs made available to drivers. All controls are software-defined, hardware enforced, and compliant with the PCI Express standards. Together, these controls allow Pensando devices to present a number of different device, management, and security profiles to the host, while maintaining strict protection against unauthorized access.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

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

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The SRK is used to encrypt (wrap) secrets (such as sensitive configuration, symmetric encryption keys, and private asymmetric keys) that can then be written to non-volatile storage, or exported off of the device for remote storage. As the SRK never leaves the RoT secure perimeter, only the origin RoT can decrypt and recover the data. The EK is used to build a chain of trust that can attest the authenticity of a DSC device. During manufacturing a Pensando-signed X.509 certificate is installed into the device which authenticates the public part of the EK. The private part of the EK never leaves the RoT. The RoT generates a 2nd-level Attestation Key-pair (authenticated via an EK-signed X.509 certificate), and that Attestation Key can be used to sign Nonce challenges. The Nonce challenge, along with the certificate chain rooted at the Pensando public Root Certificate, allows an external entity, such as Pensando's Policy and Services Manager (PSM), to authenticate the DSC as genuine. Secure Boot Secure Boot technology ensures that only software from a trusted source is allowed to run on the Pensando Distributed Services Card. At device power-on / reset, the <i>PenTrust</i> RoT processor begins executing from its immutable trusted program ROM, part of the ASIC die. The main ARM application processors are held in reset while the RoT performs its self-tests and initialization. The RoT derives its SRK from the PUF and fetches its Critical Security Parameters (CSPs) from read-only flash. These SRK-encrypted parameters include both RoT firmware and main application software authentication and decryption keys. As noted above, the use of the SRK ensures that the CSPs cannot be forged even by direct physical flash attack. The RoT loads its runtime firmware from flash into its secure, isolated memory, authenticates and decrypts it using keys from the CSP, and then passes control to it. Two Flash-based images are supported for redundancy / upgrade of the <i>PenTrust</i> firmware. The RoT then proceeds to load the main ARM application processor's first-stage boot block into secure memory that is external to the RoT (accessible to the application processor). The application processor's authentication and decryption keys are used to verify the application processor's image, and if successful the primary ARM core is taken out of reset and begins executing its trusted, authenticated code. The primary ARM application processor boot sequence follows the ARM Trusted Firmware boot flow, with a multi-step sequence bringing the system up from first-stage boot to fully operational, running Linux with an authenticated filesystem image. At each step, newly loaded programs, keys, or data are authenticated by the previous step, creating an unbroken chain-of-trust back to the RoT ROM. Dual full-system images are provided for redundancy / upgrade, and a read-only Golden Image provides a recovery option should both main images fail.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Secure Memory Regions Pensando Programmable Services Processors implement the ARM TrustZone® architecture, providing for Secure/Trusted and Normal application memory spaces. The support is implemented throughout the device, and covers hardware DMA engines, DRAM memory (HBM), and even individual memory-mapped hardware registers. Memory regions can be configured as secure, with secure memory regions being accessible only by secure hardware components. Secure components include: • ARM software running at a secure level (e.g. EL3, S-EL1) • The PenTrust RoT • The PenAccel Data Plane Crypto DMA Secure memory access is enforced in hardware. Specifically, every internal bus master has its memory transactions policed by a security filter. For memory ranges declared as secure, only certain masters are permitted access, and then only if their transactions are marked as secure (allowing a trusted master to be specific about individual transactions). The security filters themselves, along with other registers that affect security, are <i>intrinsically secure</i>; meaning they require secure access. In practice this means that overall security control is under the purview of ARM software running at a secure level. The PenAccel Crypto DMA supports encryption keys placed in secure memory, out of reach of a non-secure attacker. The high-performance encryption engines can read key material using secure memory transactions, but are restricted to using normal memory transactions to fetch data to be encrypted/decrypted. The configuration is itself secure-only, preventing a non-secure attacker from abusing the system to expose keys as data. Security Architecture of Pensando’s Distributed Services Platform, Pensando White Paper (2020), CV-AMD0017543.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Encryption Services Architecture The IPsec <u>data path</u> as well as <u>handshake</u> and <u>key management</u> 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.  The diagram illustrates the Pensando P4 Programmable Processor architecture. It features a central 'Secure Memory' block (yellow) connected by bidirectional arrows to the 'PenAccel™ Data Plane Crypto Subsystem' (left, blue) and the 'PenTrust™ Subsystem' (right, blue). Below the PenAccel subsystem is a 'Data Plane' label with a downward arrow. The entire processor is enclosed in a dashed box labeled 'Pensando P4 Programmable Processor'. Pensando IPsec Solutions, Pensando White Paper (2021), CV-AMD0016276.
[2] The programmable device of claim 1, further comprising: the router communicative to translate between a first protocol and a second protocol, for	The Accused Products are a programmable device comprising the router communicative to translate between a first protocol and a second protocol for messaging with system resources. For example, the packet buffer, P4 Ingress and P4 Egress (router) are coupled with system resources such as, for example, storage and crypto offloads, DDR, and ARM cores: The packet also has capabilities of translating between a first and second protocol. For example, the P4 Ingress and Egress pipelines can handle the new protocols as well as supporting tunnel encapsulation and decapsulation.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
messaging with system resources.	 X28794-110623 https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016111.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	• The Inline-crypto blocks (IS block) are positioned next to the P4I and P4E engines. There is one IS block after P4I (between the P4I deparser and PB), and another identical IS block after P4E (between the P4E deparser and PB). • P4I and P4E can access these engines for inline, line-rate performance encryption. The IS block supports AES-GCM 128/256 bit Key symmetric algorithms for IPsec/DTLS-like packet-based encryption use cases. In this case, the packet stays in the packet buffer while the IS block encrypts/decrypts the packet in place; therefore, no redirection/copy to memory is required. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016114. The P4 pipelines that the router is communicative to assist the packet buffer to translate between different protocols: In contrast, the AMD Pensando DPU architecture is purpose-built to accelerate services in silicon. The programmable silicon consists of <i>match processor units</i> (MPUs) capable of executing a variety of services. Using a toolchain that is centered around the P4 language, the AMD Pensando DPU can be programmed to accelerate a variety of services and protocols. Due to its programmable nature, it can also be programmed to <u>handle new protocols</u> or customized frame types that include metadata for special handling. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015806. The AMD DPU provides a programmable high-performance pipeline consisting of Rx and Tx policers, full Rx and Tx statistics (including all drop reasons), VxLAN tunnel encapsulation and decapsulation, IPv4/v6 routing, stateless and stateful security rules, network address translation (NAT), server load balancing, encryption services, VLAN to VPC mapping, and VPC peering (i.e., tunneling), IPsec, and storage offloads. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-sdn-policy-offload-user-guide.pdf, AMD-CV-0004621.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The reference pipeline supplies P4₁₆ source code for P4I and P4E modules for <u>protocol processing</u> and P4₁₆ based libraries in binary format for <u>P4 RxDMA and P4 TxDMA</u> that you can <u>access via APIs</u> for handling message transfer with host and local CPU. This combination of source code and libraries facilitates the implementation of the reference pipeline. The P4 binary can be deployed along with your P4 code on the same DPU to integrate its functionality into your system. This approach streamlines the implementation process and ensures seamless compatibility. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-sdn-policy-offload-user-guide.pdf, CV-AMD0016107. The Elba DPU has the following inline IPsec encryption/decryption features: • <u>Programmable AES-GCM inline encryption/decryption offload support in P4-Ingress and P4-Egress pipelines.</u> • <u>Supports ESP (Encapsulating Security Payload) and AH (Authentication Header) protocols.</u> • The Inline-crypto blocks (IS block) are positioned next to the P4I and P4E engines. There is one IS block after P4I (between the P4I deparser and PB), and another identical IS block after P4E (between the P4E deparser and PB). • P4I and P4E can access these engines for inline, line-rate performance encryption. The IS block supports AES-GCM 128/256 bit Key symmetric algorithms for IPsec/DTLS-like packet-based encryption use cases. In this case, <u>the packet stays in the packet buffer while the IS block encrypts/decrypts the packet in place; therefore, no redirection/copy to memory is required.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016114.
[3] The programmable device of claim 1, wherein each of the	The Accused Products include sandboxes wherein each of the sandboxes comprises user-programmable one or more of: analog circuitry, digital circuitry, memory, processors or controllers.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
sandboxes comprises user-programmable one or more of: analog circuitry, digital circuitry, memory, processors or controllers.	For example, the AMD Pensando DSC includes processors named Match Processor Units (MPU) capable of executing a various infrastructure offload service. AMD Pensando Software-in-Silicon Architecture The differentiator in performance is the architecture. A DPU that relies heavily on general-purpose cores is not the best architecture for offloading services. This class of DPU does provide the benefit of moving services off the host cores, and when combined with other accelerators they can provide fast path services, but they lack the ability to provide agility and scale for fast path services. For services that are not available in the DPU's hardware accelerators, the functionality must be executed in the general-purpose cores. In contrast, the AMD Pensando DPU architecture is purpose-built to accelerate services in silicon. <u>The programmable silicon consists of <i>match processer units</i> (MPUs) capable of executing a variety of services.</u> Using a toolchain that is centered around the P4 language, the AMD Pensando DPU can be programmed to accelerate a variety of services and protocols. Due to its programmable nature, it can also be programmed to handle new protocols or customized frame types that include metadata for special handling. AMD Pensando White Paper: Architecture Matters, Performance of a Multi-Stage SDN Pipeline on Arm vs AMD Pensando Programmable Silicon, https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015806.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Capri (16 nm) Block Diagram  https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Programmability in the DPUs DPUs have a <i>programmable data path</i>. AMD Pensando DPUs are capable of running <u>P4</u> programs natively in their <u>Match Processing Units (MPUs)</u>. P4 is a standardized open-source domain-specific language to program I/O processing devices on how to handle network traffic. P4 Program  https://community.amd.com/t5/corporate/amd-pensando-dpu-software/ba-p/630282, CV-AMD0015984.
[4] The programmable device of claim 1, wherein the port comprises: a bridge being lockable to prevent reconfiguration of the bridge and prevent reconfiguration of the router while enabling reconfiguration of	The Accused Products include the port wherein the port comprises: a bridge being lockable to prevent reconfiguration of the bridge and prevent reconfiguration of the router while enabling reconfiguration of the one or more sandboxes. <i>See claim element 1d.</i> For example, the port included in the AMD Pensando DSC's first region has a bridge that is lockable to prevent user access (including reconfiguration), and unlockable to enable user access (including reconfiguration), to each of the sandboxes.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
the one or more sandboxes.	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0016148. Based upon information and belief, the AMD Pensando DPU does not allow user to directly access the P4 engines or the P4 pipelines. To program the P4 engines, the user must log in with SSDK compiler.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Copying an Image to the DSC After building the DSC firmware image you must copy it to the DSC. You can use any of the following in-band or out-of-band management interfaces: • The internal management NIC (int_mnic) interface • The memtun interface • A serial cable connected to the DCS console • The external RJ45 out-of-band (OOB) interface Using the Internal Management NIC Interface To copy the DSC firmware image using the internal management NIC (int_mnic) interface: 1. Boot up the host with the DSC plugged in. 2. <u>Login to the host.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016120.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando you can log into the DPU, but really the way you are meant to program the part is through the P4 engines.  AMD Pensando DSC2 100 100G 2P QSFP56 DPU Login Top https://www.servethehome.com/hands-on-with-an-amd-pensando-elba-dpu-in-a-secret-lab-arm-nvidia-dell-vmware-esxi-upt/, CV-AMD0016077.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando has various hardware security features to prevent unauthorized access to the P4 engines. For example, AMD Pensando has a hardware based firewall at the PCIe interface which can restrict the user access to the P4 engines. Accordingly, it restricts any user to access the P4 engines as well as ARM cores. Pensando's Security Centric Solution Hardware-Enforced Security The highest level of security assurance is best achieved with isolated hardware that is architected to be secure from all expected threats, both internal and external. Pensando has developed the Capri Programmable Services Processor (PSP), based on a domain-specific architecture that integrates a wide range of security features described in this document. This device is deployed on Pensando's Distributed Services Cards (DSCs) that are installed in data center servers to bring secure services directly to the server edge. HardGap™ Firewall from Host In order to provide isolation from server-based vulnerabilities including deliberate attack or a rogue process writing through memory, reliable protection of host PCIe access to the DSC memory regions is required. Pensando has developed a unique hardware-based firewall at the PCIe-memory interface inside the Programmable Services Processor. Unlike traditional PCIe-attached devices, Pensando Programmable Services Processors provide no intrinsic device access to the host complex. In other words, the host has no de-facto privilege to control the device. The PCIe device tree that the host system sees is established entirely by DSC trusted software, and enforced by hardware PCIe transaction filters. The DSC software controls the PCIe devices presented by the DSC to the host, as well as the individual resource BARs made available to drivers. All controls are software-defined, hardware enforced, and compliant with the PCI Express standards. Together, these controls allow Pensando devices to present a number of different device, management, and security profiles to the host, while maintaining strict protection against unauthorized access.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

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

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The SRK is used to encrypt (wrap) secrets (such as sensitive configuration, symmetric encryption keys, and private asymmetric keys) that can then be written to non-volatile storage, or exported off of the device for remote storage. As the SRK never leaves the RoT secure perimeter, only the origin RoT can decrypt and recover the data. The EK is used to build a chain of trust that can attest the authenticity of a DSC device. During manufacturing a Pensando-signed X.509 certificate is installed into the device which authenticates the public part of the EK. The private part of the EK never leaves the RoT. The RoT generates a 2nd-level Attestation Key-pair (authenticated via an EK-signed X.509 certificate), and that Attestation Key can be used to sign Nonce challenges. The Nonce challenge, along with the certificate chain rooted at the Pensando public Root Certificate, allows an external entity, such as Pensando's Policy and Services Manager (PSM), to authenticate the DSC as genuine. Secure Boot Secure Boot technology ensures that only software from a trusted source is allowed to run on the Pensando Distributed Services Card. At device power-on / reset, the <i>PenTrust</i> RoT processor begins executing from its immutable trusted program ROM, part of the ASIC die. The main ARM application processors are held in reset while the RoT performs its self-tests and initialization. The RoT derives its SRK from the PUF and fetches its Critical Security Parameters (CSPs) from read-only flash. These SRK-encrypted parameters include both RoT firmware and main application software authentication and decryption keys. As noted above, the use of the SRK ensures that the CSPs cannot be forged even by direct physical flash attack. The RoT loads its runtime firmware from flash into its secure, isolated memory, authenticates and decrypts it using keys from the CSP, and then passes control to it. Two Flash-based images are supported for redundancy / upgrade of the <i>PenTrust</i> firmware. The RoT then proceeds to load the main ARM application processor's first-stage boot block into secure memory that is external to the RoT (accessible to the application processor). The application processor's authentication and decryption keys are used to verify the application processor's image, and if successful the primary ARM core is taken out of reset and begins executing its trusted, authenticated code. The primary ARM application processor boot sequence follows the ARM Trusted Firmware boot flow, with a multi-step sequence bringing the system up from first-stage boot to fully operational, running Linux with an authenticated filesystem image. At each step, newly loaded programs, keys, or data are authenticated by the previous step, creating an unbroken chain-of-trust back to the RoT ROM. Dual full-system images are provided for redundancy / upgrade, and a read-only Golden Image provides a recovery option should both main images fail.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Secure Memory Regions Pensando Programmable Services Processors implement the ARM TrustZone® architecture, providing for Secure/Trusted and Normal application memory spaces. The support is implemented throughout the device, and covers hardware DMA engines, DRAM memory (HBM), and even individual memory-mapped hardware registers. Memory regions can be configured as secure, with secure memory regions being accessible only by secure hardware components. Secure components include: • ARM software running at a secure level (e.g. EL3, S-EL1) • The PenTrust RoT • The PenAccel Data Plane Crypto DMA Secure memory access is enforced in hardware. Specifically, every internal bus master has its memory transactions policed by a security filter. For memory ranges declared as secure, only certain masters are permitted access, and then only if their transactions are marked as secure (allowing a trusted master to be specific about individual transactions). The security filters themselves, along with other registers that affect security, are <i>intrinsically secure</i>; meaning they require secure access. In practice this means that overall security control is under the purview of ARM software running at a secure level. The PenAccel Crypto DMA supports encryption keys placed in secure memory, out of reach of a non-secure attacker. The high-performance encryption engines can read key material using secure memory transactions, but are restricted to using normal memory transactions to fetch data to be encrypted/decrypted. The configuration is itself secure-only, preventing a non-secure attacker from abusing the system to expose keys as data. Security Architecture of Pensando’s Distributed Services Platform, Pensando White Paper (2020), CV-AMD0017543.
[5] The programmable device of claim 1, wherein the port comprises: a bridge having a security layer to authenticate a user or verify a permission, to reconfigure one or	The Accused Products have a port wherein the port comprises: a bridge that has a security layer to authenticate a user or verify a permission, to reconfigure one or more of the one or more sandboxes. For example, when users upload a new firmware image on the AMD Pensando DSC, the security layer first authenticates the user.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
more of the one or more sandboxes.	Copying an Image to the DSC After building the DSC firmware image you must copy it to the DSC. You can use any of the following in-band or out-of-band management interfaces: • The internal management NIC (int_mnic) interface • The memtun interface • A serial cable connected to the DCS console • The external RJ45 out-of-band (OOB) interface Using the Internal Management NIC Interface To copy the DSC firmware image using the internal management NIC (int_mnic) interface: 1. Boot up the host with the DSC plugged in. 2. <u>Login to the host.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf. CV-AMD0016120.
[6] The programmable device of claim 1, further comprising: the router configured to perform analytics on resources of the programmable device, on a per module basis, including determining communication bandwidth and resource utilization.	The Accused Products are programmable devices including the router being configured to perform analytics on resources of the programmable device, on a per module basis, including determining communication bandwidth and resource utilization. For example, using the “elbmon tool” a user can obtain performance analytics from the programmable device that include resource utilization, and communication bandwidth metrics.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	2.4 ASIC Performance Statistics The elbmon tool can provide a lot of useful information about the performance of the datapath. It can identify bottlenecks in the pipeline (e.g. too many instructions, too many lookups, DDR bottlenecks etc.). This is a standalone tool that provides a lot of raw information but a user friendly version of this output is provided by the following CLI. It will provide: 1. Per pipeline (P4 ingress, egress, rxdma, txdma, sxdma) and per stage MPU utilization 2. Drop count, recirc count, phv and packet counts 3. SRAM and DDR latencies 4. DDR bandwidth usage for read and write 5. Cycles per instruction 6. Cache related counters (hits vs miss) #pdsctl show system statistics asic <pre> DDR statistics: DDR Channel: 0 Number of reads: 2368732 Read latency(max/avg): 769 /107 Read TPS (in million): 0.000000 Read bandwidth: 2.353665 Number of writes: 45848 Write latency(max/avg): 601 /17 Write TPS (in million): 0.000000 Write bandwidth: 0.040825 DDR Channel: 1 Number of reads: 2358582 Read latency(max/avg): 784 /113 Read TPS (in million): 0.000000 Read bandwidth: 2.341295 Number of writes: 43669 Write latency(max/avg): 533 /15 Write TPS (in million): 0.000000 Write bandwidth: 0.036969 Pipeline statistics: Pipeline: 0 ... Pipeline: 1 PHVs: 7650 Drops: 0 Recircs: 0 Packets from PB: 7650 Parser drops: 10 Stage: 0 MPU: 0 Utilization: 0 CPI: 0.000000 MPU: 1 Utilization: 0 CPI: 0.000000 MPU: 2 Utilization: 0 CPI: 1.000000 MPU: 3 Utilization: 0 CPI: 1.000000 Stage: 1 MPU: 0 Utilization: 0 CPI: 0.000000 MPU: 1 Utilization: 0 CPI: 0.000000 MPU: 2 Utilization: 0 CPI: 1.000000 MPU: 3 Utilization: 0 CPI: 0.000000 Stage: 2 </pre> AMD-CV-0000661.
[7] The programmable device of claim 1,	The Accused Products include sandboxes, wherein the one or more sandboxes each comprise port interfaces that are configurable as to data width or speed.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
wherein the one or more sandboxes each comprise port interfaces that are configurable as to data width or speed.	For example, one or more of the sandboxes comprise a scheduler for interfacing traffic and controlling data speed: The diagram illustrates a P4 TxDMA architecture. On the left, there is a vertical stack of 'sram' and 'HBM Cache'. The main processing area is a grid of 'TE' (Transaction Engine) and 'MPU' (Match Processing Unit) blocks. A 'Scheduler' is positioned at the top right of the grid, and a 'DMA' (Direct Memory Access) block is at the bottom. A purple arrow points upwards from the Scheduler to the 'P4 TxDMA' label. A green arrow points downwards from the DMA. A red arrow points to the Scheduler. The grid consists of 8 rows and 4 columns of MPU blocks, with TE blocks to their left. https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016002.
[8] The programmable device of claim 1, wherein the one or more sandboxes each comprise clock cycle utilization regulation for controlling power consumption or	The Accused Products include sandboxes, wherein one or more sandboxes comprise clock cycle utilization regulation for controlling power consumption or thermal characteristics. For example, each Match Processing Unit (MPU) within the sandboxes is said to “operate up to 2 GHz,” while similarly each of the general-purpose ARM cores are said to “operate up to 2.8 GHz,” thus indicating that clock cycle utilization regulation takes place for the purpose of controlling power consumption and/or thermal characteristics by scaling down processing speed/frequency.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
thermal characteristics.	ASIC features • Network Interfaces: Dual 200GE, Quad 100GE, 50GE, 25GE, 10GE; also single 1GE management port • PCIe Interfaces: 32 lanes Gen4, can be divided amongst 8 ports, each port can connect to a host or root complex, max 16 lanes to one port. • P4 Pipelines: Ingress and Egress network feature pipelines each include 8 stages, each stage contains one table engine and 4 multi-issue MPUs which operate up to 2 GHz. • P4+ Pipelines: Transmit and Receive network to host/storage pipelines each include 8 stages, each stage contains one table engine and 4 MPUs which operate up to 2 GHz. New socket DMA pipelines contain 4 stages, intended for data structure management and maintenance. • P4 Large TCAM: 20 Mbits of TCAM to support ACLs and other searches from P4. • SOC: 16 ARM A-72 CPU cores operate up to 2.8 GHz. Each core has a 48 KB I-cache and 32 KB D-cache, a set of 4 cores shares a 2 MB L2 cache. Quad cores are connected via cache-coherent interconnect. Secure boot, QSPI flash, and EMMC provide SOC storage for embedded Linux. AMD-CV-0001962.
[9] The programmable device of claim 1, wherein the one or more sandboxes comprises a plurality of sandboxes that are logically or physically combinable to form a larger sandbox.	The Accused Products include sandboxes, wherein the one or more sandboxes comprises a plurality of sandboxes that are logically or physically combinable to form a larger sandbox. For example, in the accused product, there are multiple P4 engines available to perform various tasks. These multiple P4 engines can be logically combined to form a single P4 pipeline.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Besides the ready-to-deploy P4 code for the P4I pipelines, P4E pipelines, binaries, and libraries for <u>RxDMA/TxDMA/SxDMA</u>, you can build, enhance, or <u>combine</u> with your own P4 code. The following is a list of such libraries: • Flow offload (64 million sessions) with connection tracking. • VLAN and VXLAN Encap/DeCap. • Route Guard: A packet can be dropped without using flow and policy rules. • NAT: Inner IP header (Source IP, Destination IP, Source port and Destination Port). The above list gets you started and below is a list of SDN services that can be implemented using the SSDK on an AMD Pensando DPU: • Virtual Networking (Overlay) ◦ Virtual IP Spaces. ◦ Tenant aware routing (VRFs). ◦ Network address translation for inter-VPC or out-bound traffic. ◦ Capable of supporting: - Static 1:1 NAT - Many to 1 - Source - Destination - Twice NAT ◦ Ingress and egress NAT policy direction. ◦ Virtual network subnets: User defined subnets, custom IPAM I2/I3 forwarding with default gateway. ◦ Network services: NACLs, firewall, security groups, load balancing, encryption, ERSPAN (vTAP), QoS etc. ◦ VPN services: IPsec and PSP encryption services. ◦ Metering • Traffic Engineering ◦ Segment routing: SRv6 and SR-MPLS. ◦ Network service mesh: Policy driven service functions chaining. SDN Policy Offload Reference Pipeline User Guide (amd.com), CV-AMD0015821.
[10pre] A programmable	To the extent the preamble is limiting, Defendants make, use, test, sell, and offer to sell the claimed programmable device for data flow processing in a user-configurable server, including by making, using, testing,

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
device for data flow processing in a user-configurable server, comprising:	selling, and offering to sell their Distributed Service Cards (“DSCs”) and Data Processing Units (“DPUs”) including at least the Elba, Capri, Giglio, and Salina (including AMD DSC2-25G, AMD DSC2-100G, and AMD DSC2-200), as well as their DPU and DSC technology when implemented in other products, services, and offerings, for example the SmartSwitch (collectively, the “Accused Products,” “Accused DPUs,” or “Accused DSCs”). The AMD accused products infringe directly and/or indirectly (through inducing and contributing to customers use of the Accused Products) by providing a programmable device for data flow processing in a user-configurable server. AMD Pensando Accused Products (AMD Distributed Service Card (DSC) or AMD Pensando Distributed Modules (DSMs)) are Data Processing Units (DPUs) and enabled system and programmable devices for data flow processing in a user-configurable server with swappable pods or cards. Further, AMD Pensando DSCs are deployed in Server, Smart Switch, and Appliance platforms for data flow processing. The AMD Pensando Data Processing Unit (DPU) features user-configurable P4 Pipelines for dynamic data processing. The AMD Pensando Accused Products infringe directly and/or indirectly by providing a programmable device for data flow processing in a user-configurable server. For example, the AMD Pensando Distributed Service Card (DSC) is a Data Processing Units (DPUs) enabled system and is a programmable device for data flow processing in a user-configurable server. Further, AMD Pensando DSCs are deployed in Server and Appliance platforms for data flow processing. The AMD Pensando Data Processing Unit (DPU) features user-configurable P4 Pipelines for dynamic data processing.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Cloud Solution Form Factor Form Factors  DPU in Server DPU in Appliance DPU in SmartSwitch https://nowlab.cse.ohio-state.edu/static/media/workshops/presentations/exacomm23/exacomm2023_final-Kimon.pdf, CV-AMD0016028.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 AMD SECURING THE DATA CENTER COMPREHENSIVE SECURITY FROM CORE TO EDGE Distributed Security Services Extend the secure perimeter Edge and Beyond Smart Switch & Appliance Secure east-west traffic and legacy endpoints Infrastructure Advanced SmartNIC Secure every endpoint Confidential Compute Secure the application environment Root of Trust Secure the CPU at boot Compute Node Advanced Telemetry Monitor and detect threats in real time Logos: AMD EPYC, PENSANDO, ALVEO, VERSAL, elastic, FORINET, Guardicore, splunk https://www.servethehome.com/amd-data-center-solutions-at-amd-financial-analyst-day-2022/amd-fad-2022-amd-securing-the-data-center/, CV-AMD0015776.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	DSC2-25G The AMD Pensando Architecture Delivers  HARDWARE SPECIFICATIONS Form Factor: Half height, Half-length PCIe PCIe Interface: 8-lane PCIe Gen 4.0 Network Ports: 2 ports SFP56 Port Configuration: 2x10G, 2x25G Processor: Elba 7nm NETWORKING, SECURITY and STORAGE SERVICES Advanced Observability: Flow-based packet telemetry, stateful connection statistics, latency metrics, drop statistics, threshold alerting, ERSPAN (bi-directional), NetFlow/IPFIX Advanced Networking: Virtual Private Networks (network overlays), L3 ECMP, Load balancing, NAT, PAT Advanced Security: Stateful firewall with connection tracking, security groups, stateless and reflexive ACLs, VPN termination (IPsec), TLS/DTLS encryption, TLS Proxy Enhanced Storage: NVMe virtualization, NVMe-oF with RDMA or TCP transport, AES-XTS data-at-rest encryption, compression, SHA-3 deduplication, CRC64/32 acceleration AMD Pensando DSC2 25G https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015842–CV-AMD0015843.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	There is also a DSC2-100G. This is the dual QSFP28 card for 2x 100GbE. DSC2-100G The AMD Pensando Architecture Delivers  HARDWARE SPECIFICATIONS Form Factor: Full height, Half-length PCIe PCIe Interface: 16-lane PCIe Gen 4.0 Network Ports: 2 ports QSFP28 Port Configuration: 2x100G Processor: Elba 7nm PERFORMANCE AND SCALE Delivers wire-speed services on each of its QSFP28 ports, including chained services such as L4 stateful firewall + IPsec encryption + load balancing NETWORKING, SECURITY and STORAGE SERVICES Advanced Observability: Flow-based packet telemetry, stateful connection statistics, latency metrics, drop statistics, threshold alerting, ERSPAN (bi-directional), NetFlow/IPFIX Advanced Networking: Virtual Private Networks (network overlays), L3 ECMP, Load balancing, NAT, PAT Advanced Security: Stateful firewall with connection tracking, security groups, stateless and reflexive ACLs, VPN termination (IPsec), TLS/DTLS encryption, TLS Proxy Enhanced Storage: NVMe virtualization, NVMe-oF with RDMA or TCP transport, AES-XTS data-at-rest encryption, compression, SHA-3 deduplication, CRC64/32 acceleration AMD Pensando DSC2 100G https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015843.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The slide is titled "SUPERCHARGING OUR PORTFOLIO" in large white letters at the top. Below the title, there are four columns, each representing a different product category: Cloud Smart Appliance: Shows a server rack. Bullet points: "Expand Existing Cloud Footprint", "Rinse & Repeat with Other Clouds". Cloud DPU's: Shows a circuit board. Bullet points: "Double Down On Existing Deployments", "Expanding into MDC's". Enterprise DPU's: Shows a circuit board. Bullet points: "AMD the Project Monterey Platform of Choice", "AMD EPYC Solutions (Future)". Smart Switch: Shows a server rack with the Aruba logo. Bullet point: "Win More Logos". At the bottom of the slide, there are logos for Dell Technologies, Hewlett Packard Enterprise, IBM, Microsoft, Oracle, and VMware. The AMD logo is at the bottom right with the tagline "together we advance." A small number "11" is in the bottom left corner. AMD Pensando Portfolio 2022 08 30 https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015843. <i>The Pensando distributed services architecture is designed to provide programmable security, network, storage, and visibility services to a data center landscape growing in scale and complexity. This domain-specific architecture combines P4 packet processing pipelines with embedded ARM CPUs and hardware offload engines in a tightly coupled, cache coherent interconnect. Millions of flows are tracked by hardware and distributed to a combination of P4 programs, ARM processes, custom hardware engines, and PCIe based hosts and devices. There are currently 16- and 7-nm ASIC implementations of this distributed services architecture, future generations are in design.</i> Michael Galles and Francis Matus, "Pensando Distributed Services Architecture" in Theme Article: Hot Chips (2021), CV-AMD0016133.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	P4 Pipelines Pensando ASICs include multiple P4³ 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 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 Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016134.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Pensando Announces P4-programmable Platform and Joins P4 Community FEB 24, 2020  Mario Baldi ABOUT THE AUTHOR Posted by Mario Baldi on February 24, 2020 Pensando Systems recently introduced a Distributed Services Card (DSC) based on a custom, fully-programmable processor optimized to execute P4 programs. The DSC is the centerpiece of a Distributed Services Platform designed to deliver high-performance, scalable networking, storage and security services. This unique approach removes the appliance bottleneck, offers unmatched visibility on traffic, and provides significant benefits for various use cases in private, public or hybrid clouds, as well as enterprise data center networks where east-west security concerns are increasingly becoming an issue. Easily installed in standard servers, the DSC, shown in the figure below, includes specialized hardware blocks to accelerate computation intensive functions, such as compression and encryption, and relies on P4 to provide dataplane programmability at wire rate. 

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Pensando Announces P4-programmable Platform and Joins P4 Community - Open Networking Foundation by Mario Baldi (2020), CV-AMD0016130; https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc/, CV-AMD0015842. AMD Pensando DSCs are deployed in server, for example, as shown below, server platforms provided by Dell, HPE, Microsoft, Oracle, include AMD Pensando DSCs to provide data flow processing. Dell servers showcase world-record performance, scale dynamically to run highly demanding applications and workloads, and in-built features such as cyber resiliency, automation, sustainability, and global services. Dell Technologies in collaboration with VMware on their Project Monterey initiative, and by partnering with Nvidia and AMD has incorporated these technologies onto the PowerEdge server portfolio to deliver an end-to-end integrated solution to enterprises. • AMD/Pensando – There are two offerings from AMD/Pensando: ▪ AMD DSC2-25G DPU with 2 ports of 10/25G based on 2nd-generation P4-programmable SoC with 16 ARM cores. ▪ AMD DSC2-100G DPU with 2 ports of 100G based on 2nd-generation P4-programmable SoC with 16 ARM cores. https://infohub.delltechnologies.com/en-us/t/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-powerededge-servers/, CV-AMD0017558–CV-AMD0017559. • AMD Is the Clear Platform of Choice for vSphere 8: Of the three DPU vendors chosen by VMware for the vSphere Distributed Services Engine (or vSphere on DPUs), AMD is the <i>only</i> vendor with supported solutions from both vDSE launch server partners (Dell and HPE). • AMD Offers More Flexibility: AMD is the only vendor at launch with 10G/25G and 100G solutions on the market. • AMD Is the Only Proven DPU Platform: AMD is the only vendor with large scale of DPU units deployed in production with hyperscalers, public clouds, and enterprise customers. https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/white-papers/vdse-with-amd-white-paper.pdf, CV-AMD0015898.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Support from AMD Pensando Partners  AMD Pensando Distributed Services Cards are available as factory-ordered options on select HPE ProLiant, HPE Apollo, and HPE Edgeline platforms in addition to availability through HPE GreenLake, HPE's market-leading as-a-service platform. Visit HP ↗  AMD Pensando Distributed Services Cards are available as factory-ordered options on select Dell products. Visit Dell ↗ https://www.amd.com/en/support/accelerators/pensando.html, CV-AMD0015777. PENSANDO DISTRIBUTED SERVICES PLATFORM FOR HPE SERVERS https://www.hpe.com/psnow/doc/a50001239enw, CV-AMD0017537.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	“Performance, efficiency and security are integral components of competitive offerings within cloud computing and hyperconverged infrastructure,” said Forrest Norrod, senior vice president and general manager, Data Center Solutions Group, AMD. “VMware vSphere 8, running on AMD Pensando DPUs, is a key step in bringing the industry closer to composable hardware systems and truly pervasive heterogeneous computing. AMD is uniquely positioned in the industry to deliver performance, efficiency and security with a distributed services platform that is currently deployed with major customers including Goldman Sachs, Microsoft Azure, NetAPP, IBM Cloud and Oracle Cloud.” Industry-Leading AMD Pensando DPUs Enable Accelerated Data Centers with VMware vSphere 8: Advanced Micro Devices, Inc. (AMD), CV-AMD0016004. As IBM began shifting its focus to scalable cloud-native architecture in 2019, it sought out infrastructure that could scale to the cloud's needs by providing high levels of performance, security and architecture resiliency. The vendor found its answer in AMD Pensando's data processing unit (DPU). https://www.sdxcentral.com/articles/interview/why-ibm-picked-amd-pensando-dpus-for-scalable-cloud-native-architecture/2023/07/, CV-AMD0015874. Accelerating Microsoft Azure with AMD DPUs Kevin Pacella Staff  3  0  8,569 05-01-2023 10:40 AM I'm pleased to share with you that AMD Pensando DPUs are accelerating Microsoft Azure. With the recent announcement of <i>Accelerated Connections</i> available now for Public Preview, Azure customers can achieve orders of magnitude higher networking performance. This new service offering is an architectural shift that decouples the DPU adapter and SDN stack from the server. https://community.amd.com/t5/corporate/accelerating-microsoft-azure-with-amd-dpus/ba-p/604170, CV-AMD0015793.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Microsoft Azure “Pensando System’s full stack solution, complete with a P4 programmable processor and hyper scale software has enabled Microsoft Azure to rapidly augment its networking and security services improving performance, latency, and scale,” said Girish Bablani, corporate vice president, Microsoft Azure. “We have seen a 40x improvement in overall cloud based connection related performance. Pensando delivered this in less than 12 months. This partnership has enabled a transformational “limitless networking” shift that will ensure our customers have access to the highest performing products and services they expect from Microsoft Azure.” https://ir.amd.com/news-events/press-releases/detail/1057/amd-expands-data-center-solutions-capabilities-with, CV-AMD0016153. Oracle Cloud “By partnering with Pensando, Oracle Cloud Infrastructure (OCI) has developed an industry leading Software Defined Networking stack, delivering infrastructure features to Oracle Cloud customers, that have the best price to performance characteristics in the cloud industry,” said Clay Magouyrk, executive vice president, Oracle Cloud Infrastructure. “Pensando’s P4 programmable platform enables OCI to continuously innovate and deliver secure and inexpensive network services with high velocity. The ability to rapidly add new features using a programmable hardware platform will help OCI leapfrog its cloud competitors. Pensando’s engineering team, consisting of many industry experts, has collaborated very successfully with OCI engineering, and have operated at times as OCI’s extension team. OCI views Pensando as a key strategic partner in achieving our goal of being the best infrastructure platform in the cloud industry.” https://ir.amd.com/news-events/press-releases/detail/1057/amd-expands-data-center-solutions-capabilities-with, CV-AMD0016153.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Furthermore, AMD Pensando DSCs or AMD Pensando DSMs are deployed in OEM vendor Smart Appliance and SmartSwitch platforms, including the HPE Aruba CX 10000 Switch Series. SmartSwitches are fundamental building blocks of a Distributed Services Platform in a top-of-rack form factor. This new category of product can be easily integrated into any new or existing enterprise and cloud environment, regardless of whether it is public, private or hybrid. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016041. Elba DPUs are also integrated into OEM vendor Smart Appliance and Smart Switch platforms, such as the HPE Aruba CX 10000 Switch Series. See the below product data sheets for product overviews, functionality and ordering specifics. • HPE Aruba Networking CX 10000 Switch Series ○ Data Sheet ○ QuickSpecs https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pensando-elba-product-brief.pdf, CV-AMD0015774.  The switch combines the best-of-breed network operating system, Aruba AOS-CX, for data center, campus, and edge in addition to the fully programmable AMD Pensando DPU. This allows the CX 10000 to deliver stateful software-defined services inline, at scale, with wire-rate performance and orders of magnitude scale and performance improvements over traditional L2/3 switches at a fraction of their TCO. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/data-sheets/aruba-cx-10000-series-datasheet.pdf, CV-AMD0015986.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
[10a] a programmable logic device (PLD) having a first region and a second region;	The Accused Products comprise a programmable logic device (PLD) having a first region and a second region. For example, the AMD Pensando DSC contains a programmable logic device (PLD) having a first region and a second region. The AMD Pensando DSC includes hardware or software modules that are programmable using P4 pipelines to accelerate and customize cloud infrastructure services. Pensando Distributed Services Card • High-speed networking connectivity - Pensando has two 200Gbps MACs which can allow for up to 2x 200GbE networking. • High-speed packet processing with specific acceleration and often programmable logic - The Elba DPU is utilizing a P4 programmable path to align with the popular language in the networking space. • A CPU core complex - Here we have a complex again of sixteen Arm A72 core complexes • Memory controllers - Here we get dual-channel DDR4/DDR5 memory support with 8-64GB. The previous generation utilized HBM, but Pensando is using DDR now to save costs and make the design more flexible since one can use slotted memory as well. • Accelerators - There are additional accelerators cryptography, compression, and data movement, among others. This is beyond the P4 network complex • PCIe Gen4 lanes - We get a up to 32x PCIe Gen4 lanes and 8 ports (for example for 8x x4 NVMe SSDs) • Security and management features - The control plane has featured here such as a hardware root of trust. These chips also have a 1GbE NCSI interface for out-of-band management. • Runs its own OS separate from a host system - Pensando says this runs Linux with DPDK support. VMware is also prioritizing this along with the NVIDIA BlueField-2 as VMware Project Monterey. Key Focus: Networking   <i>Pensando DSC DPU Example Q2 2021</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147.
	Further, the PLD included in the Accused Products have a first region and the second region which include multiple hardware and/or software modules. For example, Coherent interconnect, P4 Ingress, P4 Egress and packet buffer are some of the modules in the first region. Similarly, P4 TxDMA, P4RxDMA, ARM cores are few examples of the modules within the second region.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  The diagram illustrates the Elba (7 nm) block architecture. At the top, it shows connections for x32 Gen4 up to 8 ports and two DDR 4/5 memory controllers, each with a 32 MB Memory Cache. The central 'Elba' block contains a PCIe MAC, VF BAR maps, a Register Interface, Interrupt Controller, eSecure, and Misc IO. It also includes SRAM, crypto, comp/decomp, CRC, and Offloads. P4 Tx/DMA and P4 Rx/DMA are connected to a central Packet Buffer Traffic Manager. The architecture features multiple P4 Egress and P4 Ingress paths, each with associated TCAM, cache, and TCAM. At the bottom, it shows Ethernet Port 0 and Ethernet Port 1, both supporting 200/100/50/40/25/10 MACs, and an NCSI interface. The diagram is credited to 'Hot Chips 32 Pensando Elba Block Diagram'. <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 FIGURE 1. Asic block diagram. Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016134. Pensando Stateful Software Services Embedded natively in the Aruba CX 10000 Switch Series, the AOS-CX operating system is an industry unique, inline stateful software service, enabled through Pensando’s <u>programmable DPU</u>, delivered at scale with wire-rate performance. The services include stateful firewalling and secure segmentation, DDoS protection, and deep flow-based session-level telemetry, enabled natively in the fabric switching infrastructure without requiring host-based agents or dedicated appliances. https://www.amd.com/system/files/documents/aruba-cx-10000-series-datasheet.pdf, AMD-CV-0004579.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Aruba CX 10000 with Pensando Help protect the unprotected with a <u>fully-programmable Aruba switching platform optimized to execute the Pensando Distributed Services Data Processing Unit (DPU) and Software Stack</u> Programmability of Pensando DSC  Figure 2: Programmability of the P4-programmable pipelines in the Pensando Distributed Services Card. Source: Principled Technologies, based on "Solution Brief: Distributed Services for Cloud Providers."⁵ Cloud providers can program the P4-programmable data plane in the Distributed Services Card. Figure 2 shows the P4-programmable data pipeline, which network engineers can leverage to customize each layer of their infrastructure stack. https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf, CV-AMD0015867.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 <i>Figure 14. Programmable pipeline</i> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015808.
[10b] the first region comprising a hardware-based or firmware-based router having a bridge with a port enabling communication between the router	The Accused Products include the first region comprising a hardware-based or firmware-based router with a port enabling communication between the router and the second region. The Accused Products include a packet buffer, P4 Ingress and P4 Egress pipelines which work together as a traffic manager and route the packets to different modules. Further, the packet buffer has multiple ports which are connected to P4 pipeline modules including P4 TxDMA, and P4RxDMA. Also, the first region has NX Coherent Interconnect which enables the communication between the first region and the second region.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
and the second region;	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible/, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0016148. Life of a Packet Packets received through the network enter on one of the Ethernet MAC interfaces and are placed in the packet buffer according to L2 or L3 class of service (COS). The packet buffer provides pause absorption buffering and COS aware arbitration for the P4 ingress pipeline. Consistent with the P4 model, packets pass through the P4 ingress pipeline for flow classification, firewall, tunnel endpoint processing, and other ingress services; return to the packet buffer in case of 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. mmi202102.issue.pdf (gibsonnet.net), CV-AMD0016199–CV-AMD00161200.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 The diagram illustrates a network processing system architecture. At the top, a 'Host Interface' is connected via a bidirectional blue arrow to a 'PCIe' block (purple). The 'PCIe' block is connected via a bidirectional purple arrow to a 'Coherent Interconnect' block (purple). The 'Coherent Interconnect' is connected via bidirectional purple arrows to three blocks: 'P4 Packet Processing Dataplane' (purple), 'ARM Cores' (purple), and 'Service Processing Offloads' (purple). These three blocks are connected via bidirectional purple arrows to a 'Packet Buffer Traffic Manager' block (green). The 'Packet Buffer Traffic Manager' is connected via bidirectional purple arrows to two 'Ethernet Port' blocks (orange). Each 'Ethernet Port' is connected via a bidirectional blue arrow to a 'Network Interface' at the bottom. The entire system is enclosed in a light blue rectangular frame. https://infohub.delltechnologies.com/en-us/l/dpus-in-the-new-vsphere-8-0-and-16th-generation-dell-powerededge-servers/amd-pensando-dpu/, CV-AMD0016074.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Figure 13. AMD Pensando Architecture A pipeline is programmed in P4 and runs in the elements that connect to the packet buffer, as depicted in Figure 14. A pipeline can include various stages that may provide one or more functions. Each stage can run in parallel on different MPUs, and as packets move through the pipeline, they are not waiting on clock cycles that may be servicing a different stage. In addition, all the MPUs have access to the same memory system, so that a packet passing from one stage to the next is not being copied to a new memory location. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015807. Further, by way of an example, it is shown below that the packet buffer is capable of routing packets to P4 pipelines such as P4 TxDMA and P4 RxDMA.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Figure: P4 Permit Action  The diagram illustrates the P4 Permit Action architecture. At the top, PCIe, DPU, and DDR are connected to a central NOC (Network On Chip). The NOC is also connected to Storage and Crypto Offloads and an Arm processor. Below the NOC are P4 TxDMA and P4 RxDMA blocks. The P4 TxDMA is connected to the P4 Egress block, and the P4 RxDMA is connected to the P4 Ingress block. Both P4 Egress and P4 Ingress blocks are connected to a central Packet Buffer Traffic Manager. The Packet Buffer Traffic Manager is connected to Eth 1/1 and Eth 1/2 blocks. Eth 1/1 is labeled 'Encrypted' and Eth 1/2 is labeled 'Unencrypted'. There are also Inline Crypto blocks connected to the Packet Buffer Traffic Manager and the P4 Ingress block. A green line traces a path from the P4 Ingress block through the Packet Buffer Traffic Manager to Eth 1/2. A small identifier 'X28794-110623' is located at the bottom right of the diagram. https://docs.amd.com/r/en-US/ug1669-amd-pensando-ssdk-hello-world-user-guide/Redirect-P4-Action, CV-AMD0016083.
[10c] the second region comprising a plurality of sandboxes with user-definable	The Accused Products comprise the second region comprising a plurality of sandboxes with user-definable programmable electronic circuits of the PLD.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
programmable electronic circuits of the PLD; and	For example, the second region comprises one or more sandboxes with user-definable programmable electronic circuits of the PLD. For example, the second region contains the P4 pipelines (P4 TxDMA, and P4RxDMA) which function as sandboxes being that users can program them to function as needed. In addition, the accused product also includes ARM cores (second region) that users can program. Capri (16 nm) Block Diagram  The diagram illustrates the Capri (16 nm) block architecture. At the top, it shows a connection to 'x16 Gen3/4' via a 'PCIe MAC' and 'VF BAR maps'. The central 'Capri' block is connected to 'HBM' (High Bandwidth Memory) and an 'NX Coherent Interconnect + 1MB HBM Cache'. Below this, there are several processing blocks: 'P4 TxDMA' and 'P4 RxDMA' (both containing 'sram', 'HBM Cache', and 'DMA' components), 'P4 Egress' (containing 'sram', 'cache', and 'TCAM'), and 'P4 Ingress' (containing 'sram', 'cache', and 'TCAM'). These are all connected to a central 'Packet Buffer Traffic Manager'. To the left, there are 'Register Interface', 'Interrupt Controller', 'timers doorbells', and 'Misc IO' blocks. To the right, there is an 'L2 Cache' and 'ARM A72' cores. At the bottom, the architecture connects to 'Ethernet Port 0' and 'Ethernet Port 1' via 'timestamp' and 'timestamp MAC' blocks, and an 'NCSI' block. https://www.servethehome.com/amd-pensando-giglio-dpu-for-2023-salina-dpu-in-2024-and-amd-epyc-CV-AMD0015842.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Programmability of Pensando DSC  Figure 2: Programmability of the P4-programmable pipelines in the Pensando Distributed Services Card. Source: Principled Technologies, based on "Solution Brief: Distributed Services for Cloud Providers."⁵ Cloud providers can program the P4-programmable data plane in the Distributed Services Card. Figure 2 shows the P4-programmable data pipeline, which network engineers can leverage to customize each layer of their infrastructure stack. https://www.principledtechnologies.com/Pensando/DSC-200-performance-0122.pdf, CV-AMD0015867.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Figure 14. Programmable pipeline In contrast, the AMD Pensando DPU architecture is purpose-built to accelerate services in silicon. The programmable silicon consists of <i>match processor units</i> (MPUs) capable of executing a variety of services. Using a toolchain that is centered around the P4 language, the AMD Pensando DPU can be programmed to accelerate a variety of services and protocols. Due to its programmable nature, it can also be programmed to handle new protocols or customized frame types that include metadata for special handling. The general-purpose Arm cores in the AMD Pensando architecture are used primarily for control plane and exception packets. Figure 13 shows the architecture of the AMD Pensando DPU. The key central components of the DPU are the MPUs; this is where the software-in-silicon is executed and provides accelerated (fast path) services. Note that the system memory is connected to both the general-purpose cores and the domain-specific MPUs. This allows features that are typically handled by general-purpose cores, like connection tracking, to be implemented on the MPUs and executed directly in the pipeline. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015806.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AGILE PLATFORM FOR CLOUD PROVIDERS The DSC-100 is the ideal software-defined platform to bring high-performance and efficiency to the cloud infrastructure. Developers can create customized data plane services that are targeted to 100G throughput, microsecond-level latencies, and that scale to millions of flows. High-level programming languages (P4, C) enable rapid development and deployment of new features and services. The programmability and agility of the DSC platform allows customized functionality to be developed by cloud providers, speeding innovation and competitive differentiation. • Incorporates both data plane and control plane, eliminating host agents • Customizable data plane - customers can develop their own software on the platform INDUSTRY-LEADING HARDWARE PLATFORM Pensando's P4 Programmable Processor powers the DSC, enabling wire-speed performance – even with chained services – as well as enhanced security through isolated enforcement of policy. The data plane and control planes are fully software-defined and supported with optimized hardware accelerators. Up to 8 GBytes of on-board DRAM provides the performance and flow-table capacity for true cloud-scale deployments. https://cmssc28productionstorage.blob.core.windows.net/article-assets/5eecb987d386d2c61a5e67b0-2.pdf, CV-AMD0016072.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement  Pensando Distributed Services Card • High-speed networking connectivity - Pensando has two 200Gbps MACs which can allow for up to 2x 200GbE networking. • High-speed packet processing with specific acceleration and often programmable logic - The Elba DPU is utilizing a P4 programmable path to align with the popular language in the networking space. • A CPU core complex - Here we a complex again of sixteen Arm A72 core complexes • Memory controllers - Here we get dual-channel DDR4/ DDR5 memory support with 8-64GB. The previous generation utilized HBM, but Pensando is using DDR now to save costs and make the design more flexible since one can use slotted memory as well. • Accelerators - There are additional accelerators cryptography, compression, and data movement, among others. This is beyond the P4 network complex • PCIe Gen4 lanes - We get a up to 32x PCIe Gen4 lanes and 8 ports (for example for 8x x4 NVMe SSDs) • Security and management features - The control plane has featured here such as a hardware root of trust. These chips also have a 1GbE NCSI interface for out-of-band management. • Runs its own OS separate from a host system - Pensando says this runs Linux with DPDK support. VMware is also prioritizing this along with the NVIDIA BlueField-2 as VMware Project Monterey. Key Focus: Networking   https://www.servethehome.com/dpu-vs-smartnic-sth-nic-continuum-framework-for-discussing-nic-types/pensando-dsc-dpu-example-q2-2021/, CV-AMD0016364.
[10d] an internal communication network in the PLD, coupling the port of the bridge of the router and the one or more sandboxes	The Accused Products comprise an internal communication network in the PLD, coupling the port of the bridge of the router and the one or more sandboxes for the data flow processing. For example, the port included in the AMD Pensando DSC's first region has a bridge that is lockable to prevent user access, and unlockable to enable user access, to each of the programmable P4 engines.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
for the data flow processing,	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900. Based upon information and belief, AMD Pensando DPU does not allow user to directly access the P4 engines or the P4 pipelines. In order to program the P4 engines, the user has to log in with SSDK compiler.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Copying an Image to the DSC After building the DSC firmware image you must copy it to the DSC. You can use any of the following in-band or out-of-band management interfaces: • The internal management NIC (int_mnic) interface • The memtun interface • A serial cable connected to the DCS console • The external RJ45 out-of-band (OOB) interface Using the Internal Management NIC Interface To copy the DSC firmware image using the internal management NIC (int_mnic) interface: 1. Boot up the host with the DSC plugged in. 2. <u>Login to the host.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016120.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando you can log into the DPU, but really the way you are meant to program the part is through the P4 engines.  AMD Pensando DSC2 100 100G 2P QSP56 DPU Login Top https://www.servethehome.com/hands-on-with-an-amd-pensando-elba-dpu-in-a-secret-lab-arm-nvidia-dell-vmware-esxi-upt/, CV-AMD016077.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Furthermore, AMD Pensando has various hardware security features to prevent unauthorized access to the P4 engines. For example, AMD Pensando has a hardware based firewall at the PCIe interface which can restrict the user access to the P4 engines. Accordingly, it restricts any user to access the P4 engines as well as ARM cores. Pensando's Security Centric Solution Hardware-Enforced Security The highest level of security assurance is best achieved with isolated hardware that is architected to be secure from all expected threats, both internal and external. Pensando has developed the Capri Programmable Services Processor (PSP), based on a domain-specific architecture that integrates a wide range of security features described in this document. This device is deployed on Pensando's Distributed Services Cards (DSCs) that are installed in data center servers to bring secure services directly to the server edge. HardGap™ Firewall from Host In order to provide isolation from server-based vulnerabilities including deliberate attack or a rogue process writing through memory, reliable protection of host PCIe access to the DSC memory regions is required. Pensando has developed a unique hardware-based firewall at the PCIe-memory interface inside the Programmable Services Processor. Unlike traditional PCIe-attached devices, Pensando Programmable Services Processors provide no intrinsic device access to the host complex. In other words, the host has no de-facto privilege to control the device. The PCIe device tree that the host system sees is established entirely by DSC trusted software, and enforced by hardware PCIe transaction filters. The DSC software controls the PCIe devices presented by the DSC to the host, as well as the individual resource BARs made available to drivers. All controls are software-defined, hardware enforced, and compliant with the PCI Express standards. Together, these controls allow Pensando devices to present a number of different device, management, and security profiles to the host, while maintaining strict protection against unauthorized access.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

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

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The SRK is used to encrypt (wrap) secrets (such as sensitive configuration, symmetric encryption keys, and private asymmetric keys) that can then be written to non-volatile storage, or exported off of the device for remote storage. As the SRK never leaves the RoT secure perimeter, only the origin RoT can decrypt and recover the data. The EK is used to build a chain of trust that can attest the authenticity of a DSC device. During manufacturing a Pensando-signed X.509 certificate is installed into the device which authenticates the public part of the EK. The private part of the EK never leaves the RoT. The RoT generates a 2nd-level Attestation Key-pair (authenticated via an EK-signed X.509 certificate), and that Attestation Key can be used to sign Nonce challenges. The Nonce challenge, along with the certificate chain rooted at the Pensando public Root Certificate, allows an external entity, such as Pensando's Policy and Services Manager (PSM), to authenticate the DSC as genuine. Secure Boot Secure Boot technology ensures that only software from a trusted source is allowed to run on the Pensando Distributed Services Card. At device power-on / reset, the PenTrust RoT processor begins executing from its immutable trusted program ROM, part of the ASIC die. The main ARM application processors are held in reset while the RoT performs its self-tests and initialization. The RoT derives its SRK from the PUF and fetches its Critical Security Parameters (CSPs) from read-only flash. These SRK-encrypted parameters include both RoT firmware and main application software authentication and decryption keys. As noted above, the use of the SRK ensures that the CSPs cannot be forged even by direct physical flash attack. The RoT loads its runtime firmware from flash into its secure, isolated memory, authenticates and decrypts it using keys from the CSP, and then passes control to it. Two Flash-based images are supported for redundancy / upgrade of the PenTrust firmware. The RoT then proceeds to load the main ARM application processor's first-stage boot block into secure memory that is external to the RoT (accessible to the application processor). The application processor's authentication and decryption keys are used to verify the application processor's image, and if successful the primary ARM core is taken out of reset and begins executing its trusted, authenticated code. The primary ARM application processor boot sequence follows the ARM Trusted Firmware boot flow, with a multi-step sequence bringing the system up from first-stage boot to fully operational, running Linux with an authenticated filesystem image. At each step, newly loaded programs, keys, or data are authenticated by the previous step, creating an unbroken chain-of-trust back to the RoT ROM. Dual full-system images are provided for redundancy / upgrade, and a read-only Golden Image provides a recovery option should both main images fail.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Secure Memory Regions Pensando Programmable Services Processors implement the ARM TrustZone® architecture, providing for Secure/Trusted and Normal application memory spaces. The support is implemented throughout the device, and covers hardware DMA engines, DRAM memory (HBM), and even individual memory-mapped hardware registers. Memory regions can be configured as secure, with secure memory regions being accessible only by secure hardware components. Secure components include: • ARM software running at a secure level (e.g. EL3, S-EL1) • The PenTrust RoT • The PenAccel Data Plane Crypto DMA Secure memory access is enforced in hardware. Specifically, every internal bus master has its memory transactions policed by a security filter. For memory ranges declared as secure, only certain masters are permitted access, and then only if their transactions are marked as secure (allowing a trusted master to be specific about individual transactions). The security filters themselves, along with other registers that affect security, are <i>intrinsically secure</i>; meaning they require secure access. In practice this means that overall security control is under the purview of ARM software running at a secure level. The PenAccel Crypto DMA supports encryption keys placed in secure memory, out of reach of a non-secure attacker. The high-performance encryption engines can read key material using secure memory transactions, but are restricted to using normal memory transactions to fetch data to be encrypted/decrypted. The configuration is itself secure-only, preventing a non-secure attacker from abusing the system to expose keys as data. Security Architecture of Pensando’s Distributed Services Platform, Pensando White Paper (2020), CV-AMD0017543.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Encryption Services Architecture The IPsec <u>data path</u> as well as <u>handshake</u> and <u>key management</u> 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.  The diagram illustrates the Pensando P4 Programmable Processor architecture. It features a central 'Secure Memory' block (yellow) connected to two subsystems: 'PenAccel™ Data Plane Crypto Subsystem' (left, blue) and 'PenTrust™ Subsystem' (right, blue). Both subsystems are connected to a 'Data Plane' at the bottom. The entire processor is enclosed in a dashed box labeled 'Pensando P4 Programmable Processor'. Pensando IPsec Solutions, Pensando White Paper (2021), CV-AMD0016276.
[10e] wherein the bridge is lockable to prevent user access, and unlockable to enable user access, to the plurality of sandboxes and user-definable programmable electronic circuits therein.	The Accused Products comprise a bridge, wherein bridge is lockable to prevent user access, and unlockable to enable user access, to the plurality of sandboxes and user-definable programmable electronic circuits therein. For example, the port included in the AMD Pensando DSC's first region has a bridge that is lockable to prevent user access, and unlockable to enable user access, to each of the programmable P4 engines.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Capri (16 nm) Block Diagram  <i>Hot Chips 32 Pensando Capri Block Diagram</i> https://www.servethehome.com/amd-acquires-pensando-for-its-dpu-future-intel-ipu-nvidia-bluefield-marvell-octeon-fungible, CV-AMD0016147–CV-AMD0016148.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Elba (7 nm) Block Diagram  <i>Hot Chips 32 Pensando Elba Block Diagram</i> https://www.servethehome.com/pensando-distributed-services-architecture-smartnic/hot-chips-32-pensando-elba-block-diagram/, CV-AMD0015900. Based upon information and belief, the AMD Pensando DPU does not allow user to directly access the P4 engines or the P4 pipelines. To program the P4 engines, the user must log in with SSDK compiler.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Copying an Image to the DSC After building the DSC firmware image you must copy it to the DSC. You can use any of the following in-band or out-of-band management interfaces: • The internal management NIC (int_mnic) interface • The memtun interface • A serial cable connected to the DCS console • The external RJ45 out-of-band (OOB) interface Using the Internal Management NIC Interface To copy the DSC firmware image using the internal management NIC (int_mnic) interface: 1. Boot up the host with the DSC plugged in. 2. <u>Login to the host.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016120.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando you can log into the DPU, but really the way you are meant to program the part is through the P4 engines.  AMD Pensando DSC2 100 100G 2P QSP56 DPU Login Top https://www.servethehome.com/hands-on-with-an-amd-pensando-elba-dpu-in-a-secret-lab-arm-nvidia-dell-vmware-esxi-upt/, CV-AMD0016077.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	AMD Pensando has various hardware security features to prevent unauthorized access to the P4 engines. For example, AMD Pensando has a hardware based firewall at the PCIe interface which can restrict the user access to the P4 engines. Accordingly, it restricts any user to access the P4 engines as well as ARM cores. Pensando's Security Centric Solution Hardware-Enforced Security The highest level of security assurance is best achieved with isolated hardware that is architected to be secure from all expected threats, both internal and external. Pensando has developed the Capri Programmable Services Processor (PSP), based on a domain-specific architecture that integrates a wide range of security features described in this document. This device is deployed on Pensando's Distributed Services Cards (DSCs) that are installed in data center servers to bring secure services directly to the server edge. HardGap™ Firewall from Host In order to provide isolation from server-based vulnerabilities including deliberate attack or a rogue process writing through memory, reliable protection of host PCIe access to the DSC memory regions is required. Pensando has developed a unique hardware-based firewall at the PCIe-memory interface inside the Programmable Services Processor. Unlike traditional PCIe-attached devices, Pensando Programmable Services Processors provide no intrinsic device access to the host complex. In other words, the host has no de-facto privilege to control the device. The PCIe device tree that the host system sees is established entirely by DSC trusted software, and enforced by hardware PCIe transaction filters. The DSC software controls the PCIe devices presented by the DSC to the host, as well as the individual resource BARs made available to drivers. All controls are software-defined, hardware enforced, and compliant with the PCI Express standards. Together, these controls allow Pensando devices to present a number of different device, management, and security profiles to the host, while maintaining strict protection against unauthorized access.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

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

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The SRK is used to encrypt (wrap) secrets (such as sensitive configuration, symmetric encryption keys, and private asymmetric keys) that can then be written to non-volatile storage, or exported off of the device for remote storage. As the SRK never leaves the RoT secure perimeter, only the origin RoT can decrypt and recover the data. The EK is used to build a chain of trust that can attest the authenticity of a DSC device. During manufacturing a Pensando-signed X.509 certificate is installed into the device which authenticates the public part of the EK. The private part of the EK never leaves the RoT. The RoT generates a 2nd-level Attestation Key-pair (authenticated via an EK-signed X.509 certificate), and that Attestation Key can be used to sign Nonce challenges. The Nonce challenge, along with the certificate chain rooted at the Pensando public Root Certificate, allows an external entity, such as Pensando's Policy and Services Manager (PSM), to authenticate the DSC as genuine. Secure Boot Secure Boot technology ensures that only software from a trusted source is allowed to run on the Pensando Distributed Services Card. At device power-on / reset, the <i>PenTrust</i> RoT processor begins executing from its immutable trusted program ROM, part of the ASIC die. The main ARM application processors are held in reset while the RoT performs its self-tests and initialization. The RoT derives its SRK from the PUF and fetches its Critical Security Parameters (CSPs) from read-only flash. These SRK-encrypted parameters include both RoT firmware and main application software authentication and decryption keys. As noted above, the use of the SRK ensures that the CSPs cannot be forged even by direct physical flash attack. The RoT loads its runtime firmware from flash into its secure, isolated memory, authenticates and decrypts it using keys from the CSP, and then passes control to it. Two Flash-based images are supported for redundancy / upgrade of the <i>PenTrust</i> firmware. The RoT then proceeds to load the main ARM application processor's first-stage boot block into secure memory that is external to the RoT (accessible to the application processor). The application processor's authentication and decryption keys are used to verify the application processor's image, and if successful the primary ARM core is taken out of reset and begins executing its trusted, authenticated code. The primary ARM application processor boot sequence follows the ARM Trusted Firmware boot flow, with a multi-step sequence bringing the system up from first-stage boot to fully operational, running Linux with an authenticated filesystem image. At each step, newly loaded programs, keys, or data are authenticated by the previous step, creating an unbroken chain-of-trust back to the RoT ROM. Dual full-system images are provided for redundancy / upgrade, and a read-only Golden Image provides a recovery option should both main images fail.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Secure Memory Regions Pensando Programmable Services Processors implement the ARM TrustZone® architecture, providing for Secure/Trusted and Normal application memory spaces. The support is implemented throughout the device, and covers hardware DMA engines, DRAM memory (HBM), and even individual memory-mapped hardware registers. Memory regions can be configured as secure, with secure memory regions being accessible only by secure hardware components. Secure components include: • ARM software running at a secure level (e.g. EL3, S-EL1) • The PenTrust RoT • The PenAccel Data Plane Crypto DMA Secure memory access is enforced in hardware. Specifically, every internal bus master has its memory transactions policed by a security filter. For memory ranges declared as secure, only certain masters are permitted access, and then only if their transactions are marked as secure (allowing a trusted master to be specific about individual transactions). The security filters themselves, along with other registers that affect security, are <i>intrinsically secure</i>; meaning they require secure access. In practice this means that overall security control is under the purview of ARM software running at a secure level. The PenAccel Crypto DMA supports encryption keys placed in secure memory, out of reach of a non-secure attacker. The high-performance encryption engines can read key material using secure memory transactions, but are restricted to using normal memory transactions to fetch data to be encrypted/decrypted. The configuration is itself secure-only, preventing a non-secure attacker from abusing the system to expose keys as data. Security Architecture of Pensando’s Distributed Services Platform, Pensando White Paper (2020), CV-AMD0017543.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	Encryption Services Architecture The IPsec <u>data path</u> as well as <u>handshake</u> and <u>key management</u> 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 <i>PenTrust</i> 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.  The diagram illustrates the Pensando P4 Programmable Processor architecture. It features three main components within a dashed box: the PenAccel™ Data Plane Crypto Subsystem (represented by a blue padlock icon), Secure Memory (represented by a yellow box), and the PenTrust™ Subsystem (represented by a blue padlock icon). Arrows indicate bidirectional communication between these components. Below the dashed box, a 'Data Plane' label is connected to the PenAccel subsystem by a vertical arrow. Pensando IPsec Solutions, Pensando White Paper (2021), CV-AMD0016276.
[11] The programmable device of claim 10, further comprising: the router to translate between a first protocol used external to the second region and a second protocol used in the internal communication network of the PLD.	The Accused Products include a programmable device further comprising the router to translate between a first protocol used external to the second region and a second protocol used in the internal communication network of the PLD. For example, the router utilizes its P4 pipelines that it is coupled to in order to translate between a first protocol used external to the second region.

Exhibit D: Claim Chart for U.S. Patent 10,985,943


Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
used in the internal communication network in the PLD.	 X28794-110623 https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016111.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	• The Inline-crypto blocks (IS block) are positioned next to the P4I and P4E engines. There is one IS block after P4I (between the P4I deparser and PB), and another identical IS block after P4E (between the P4E deparser and PB). • P4I and P4E can access these engines for inline, line-rate performance encryption. The IS block supports AES-GCM 128/256 bit Key symmetric algorithms for IPsec/DTLS-like packet-based encryption use cases. In this case, the packet stays in the packet buffer while the IS block encrypts/decrypts the packet in place; therefore, no redirection/copy to memory is required. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016114. For example, the router may take in encrypted packets, and decrypt them for use within the second region and internal communication network: In contrast, the AMD Pensando DPU architecture is purpose-built to accelerate services in silicon. The programmable silicon consists of <i>match processer units</i> (MPUs) capable of executing a variety of services. Using a toolchain that is centered around the P4 language, the AMD Pensando DPU can be programmed to accelerate a variety of services and protocols. Due to its programmable nature, it can also be programmed to <u>handle new protocols</u> or customized frame types that include metadata for special handling. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-comparison-of-dpu-hardware-strategies.pdf, CV-AMD0015806. The AMD DPU provides a programmable high-performance pipeline consisting of Rx and Tx policers, full Rx and Tx statistics (including all drop reasons), VxLAN tunnel encapsulation and decapsulation, IPv4/v6 routing, stateless and stateful security rules, network address translation (NAT), server load balancing, encryption services, VLAN to VPC mapping, and VPC peering (i.e., tunneling), IPsec, and storage offloads. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-sdn-policy-offload-user-guide.pdf, AMD-CV-0004621.

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	The reference pipeline supplies P4₁₆ source code for P4I and P4E modules for <u>protocol processing</u> and P4₁₆ based libraries in binary format for <u>P4 RxDMA and P4 TxDMA</u> that you can <u>access via APIs</u> for handling message transfer with host and local CPU. This combination of source code and libraries facilitates the implementation of the reference pipeline. The P4 binary can be deployed along with your P4 code on the same DPU to integrate its functionality into your system. This approach streamlines the implementation process and ensures seamless compatibility. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-sdn-policy-offload-user-guide.pdf, CV-AMD0015817. The Elba DPU has the following inline IPsec encryption/decryption features: • <u>Programmable AES-GCM inline encryption/decryption offload support in P4-Ingress and P4-Egress pipelines.</u> • <u>Supports ESP (Encapsulating Security Payload) and AH (Authentication Header) protocols.</u> • The Inline-crypto blocks (IS block) are positioned next to the P4I and P4E engines. There is one IS block after P4I (between the P4I deparser and PB), and another identical IS block after P4E (between the P4E deparser and PB). • P4I and P4E can access these engines for inline, line-rate performance encryption. The IS block supports AES-GCM 128/256 bit Key symmetric algorithms for IPsec/DTLS-like packet-based encryption use cases. In this case, <u>the packet stays in the packet buffer while the IS block encrypts/decrypts the packet in place; therefore, no redirection/copy to memory is required.</u> https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/user-guides/pensando-ssdk-ipsec-gw-user-guide.pdf, CV-AMD0016114.
[12] The programmable device of claim 10, wherein each of the	The Accused Products include sandboxes wherein each of the plurality of sandboxes comprises user-programmable one or more of: analog circuitry, digital circuitry, memory, processors or controllers. <i>See claim 3.</i>

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
plurality of sandboxes comprises user-programmable one or more of: analog circuitry, digital circuitry, memory, processors or controllers.	
[13] The programmable device of claim 10, wherein the bridge is lockable to enable reconfiguration of the plurality of sandboxes, prevent reconfiguration of the bridge and prevent reconfiguration of the router.	The Accused Products comprise a bridge, wherein the bridge is lockable to enable reconfiguration of the plurality of sandboxes, prevent reconfiguration of the bridge and prevent reconfiguration of the router. <i>See claim element 10e.</i>
[14] The programmable device of claim 10, wherein the bridge has a security layer to authenticate a user or verify a permission, to	The Accused Products comprise a bridge, wherein the bridge has a security layer to authenticate a user or verify a permission, to reconfigure one or more of the plurality of sandboxes. <i>See claim 5.</i>

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
reconfigure one or more of the plurality of sandboxes.	
[15] The programmable device of claim 10, further comprising: the router configured to perform analytics on resources of the programmable device, on a per module basis, including determining communication bandwidth and resource utilization in the plurality of sandboxes.	The Accused Products are programmable devices including a router configured to perform analytics on resources of the programmable device, on a per module basis, including determining communication bandwidth and resource utilization in the plurality of sandboxes. <i>See claim 6.</i>
[16] The programmable device of claim 10, wherein the each of the plurality of sandboxes comprises one or more port interfaces configurable as to data width or speed.	The Accused Products include sandboxes, wherein each of the plurality of sandboxes comprises one or more port interfaces configurable as to data width or speed. <i>See claim 7.</i>

Exhibit D: Claim Chart for U.S. Patent 10,985,943

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
[17] The programmable device of claim 10, wherein each of the plurality of more sandboxes comprises clock cycle utilization regulation for controlling power consumption or thermal characteristics.	The Accused Products include sandboxes, wherein each of the plurality of more sandboxes comprises clock cycle utilization regulation for controlling power consumption or thermal characteristics. <i>See claim 8.</i>
[18] The programmable device of claim 10, wherein the plurality of sandboxes is logically or physically combinable to form a single, larger sandbox.	The Accused Products include sandboxes, wherein the plurality of sandboxes is logically or physically combinable to form a single, larger sandbox. <i>See claim 9.</i>