

Exhibit E: Claim Chart for U.S. Patent 10,944,634

U.S. Patent No. 10,944,634

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. 10,944,634 (the “’634 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 ’634 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, including as described herein in the networks in which the patented technology is implemented. 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 networks in which the patented technology is implemented and are 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 in the infringing networks. 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. Defendants and their customer’s and client’s networks that implement or include the Accused Products are the accused networks. Further details are provided herein.

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 E: Claim Chart for U.S. Patent 10,944,634

Theories exemplifying how the Accused Products infringe the asserted claims of the '634 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 of the examples provided, and include generally the same components. For example, for purposes of infringement of the '634 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' Responses 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' Responses 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 method for network connection optimization, comprising:	To the extent the preamble is limiting, Defendants perform the claimed method for network connection optimization, including by making, using, testing, selling, and offering to sell the Accused Products; implementing and/or testing those Accused Products in their own networks; and providing them to customers. <i>See, e.g.</i>, AMD-CV-0003945–AMD-CV-0004184. AMD Pensando Distributed Services Cards (DSCs) or AMD Pensando Distributed Modules (DSMs) are Data Processing Unit (DPU) enabled systems deployed in Server, Smart Switch, and Appliance platforms for data flow processing and for providing various infrastructure services. Further, the AMD Pensando DPUs deployed in the AMD Pensando DSCs or SmartSwitch Solutions function as a layer 3 switch and provide further stateful services such as firewall, load balancer, and telemetry. It includes flow-based engine in hardware that implements a packet forwarding mechanism.

Exhibit E: Claim Chart for U.S. Patent 10,944,634

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	2020: SmartSwitches (Distributed Service Switches) SmartSwitches, a.k.a Distributed Service Switches, add to the state of the art layer 3 switches a rich collection of wire-rate stateful services, e.g., firewall, LB (Load Balancer), NAT (Network Address Translation, <u>TAP (Test Access Point) Network (traffic observability)</u>, streaming telemetry, etc. A SmartSwitch can apply these services both to the underlay network and to the overlay network. In the second case, they are typically colocated with the VTEPs. To effectively support these services, SmartSwitches have a different forwarding mechanism that is stateful: decision of what to do with the packet may depend on previous packets of the same communication. Hence, each packet between host A and host B is processed as a part of a session between A and B that comprises two unidirectional flows: one from A to B and one from B to A. The term session applies as well to connection oriented protocols like TCP or connectionless protocols like UDP. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016048. The Pensando flow-based engine stores each flow in a flow table. Each flow table entry includes counters and timers. For the first time in the history of networking, <u>it is possible to provide precise information on how each flow behaves, including parameters like bytes and packet exchange, delays, jitter, and bandwidth.</u> Previously this information was available only at the aggregate level, for example, at the interface level. This new granularity is what we call "Telemetry," It allows us to pinpoint performance issues at the network and application level. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016066.

Exhibit E: Claim Chart for U.S. Patent 10,944,634

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
sending the parameter values for determination of bandwidth; from the first endpoint device, to the tuner server that is configured to receive further parameter values for determination of bandwidth of network connections from a plurality of endpoint devices;	For example, the flow-based engine stored within the AMD Pensando DPU deployed in DSC or SmartSwitch stores the packets for network connection sessions between two endpoint devices as flows in a flow table. Using the five fields of a flow (IP source and destination addresses, protocol type, TCP/UDP source destination ports and other information like input interface, MAC addresses, TCP flags, etc.), the AMD Pensando flow-based engine can extract parameters such as bytes and packet exchange, delays, jitter and bandwidth. On information and belief, these parameter values are sent from the first endpoint device to the tuner server. Further, the tuner server is configured to receive each of the parameters values, “further parameter values.” See claim element 1c. The Key Difference in Architecture: The Pensando Flow-Based Engine We have previously mentioned that: <i>SmartSwitches have a different forwarding mechanism that is stateful. Each packet between host A and host B is processed as a part of a session between A and B. The session comprises two unidirectional flows, one from A to B and one from B to A.</i> Sessions and flows are identified using an extended five-tuple that includes the basic five fields (IP source and destination addresses, protocol type, TCP/UDP source destination ports) and other information like input interface, MAC addresses, TCP flags, etc. When you build hardware capable of implementing this new forwarding mechanism at wire speed, many new features become possible. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016065.

Exhibit E: Claim Chart for U.S. Patent 10,944,634

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. Michael Galles and Francis Matus, “Pensando Distributed Services Architecture” in Theme Article: Hot Chips (2021), CV-AMD0016133–CV-AMD0016134.
[2c] receiving, at the first endpoint device, from the tuner server, the one	The AMD Pensando Accused Products (including the networks in which they are implemented) receive, at the first endpoint device, from the tuner server, the one or more para values for the transmission bandwidth. For example, once the flow-based engine gathers the transmission bandwidth for the past network connection (e.g., a flow in the flow table) based on the five fields (e.g., parameter values) of a flow from, on information and belief, the


Exhibit E: Claim Chart for U.S. Patent 10,944,634

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
or more parameter values for the transmission bandwidth.	first endpoint device, the flow-based engine forwards the relevant packet to initiate the next network connection from the first endpoint device (e.g., a server) to the second endpoint device (e.g., another server), based on the determination at the AMD Pensando DPU. In this process, on information and belief, the first endpoint device receives the parameter values for the transmission bandwidth. <i>See claim element 1c.</i> The Key Difference in Architecture: The Pensando Flow-Based Engine We have previously mentioned that: <i>SmartSwitches have a different forwarding mechanism that is stateful. Each packet between host A and host B is processed as a part of a session between A and B. The session comprises two unidirectional flows, one from A to B and one from B to A.</i> Sessions and flows are identified using an extended five-tuple that includes the basic five fields (IP source and destination addresses, protocol type, TCP/UDP source destination ports) and other information like input interface, MAC addresses, TCP flags, etc. When you build hardware capable of implementing this new forwarding mechanism at wire speed, many new features become possible. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/white-papers/pensando-smartswitches.pdf, CV-AMD0016065.

Exhibit E: Claim Chart for U.S. Patent 10,944,634

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement																				
[5b] adjusting the transmission bandwidth based on the monitoring.	The AMD Pensando Accused Products (including the networks in which they are implemented) adjust the transmission bandwidth based on the monitoring. For example, in addition to the telemetry services, AMD Pensando DPUs also act as load balancer and thus, based on the monitored network traffic, it adjusts the transmission bandwidth between the first endpoint device and the second endpoint device. Data Processing Units (DPUs) have become the backbone of next-generation data centers, powering many accelerated network functions (NF). AMD Pensando™ programmable DPUs provide many functions, such as: <table><tr><th>Software-Defined Networking</th><th colspan="2">Software-Defined Storage</th><th colspan="2">PCIe® Services</th></tr><tr><td> • Policy management • Virtual networking • NFV • vSwitch • NAT / load balancer • Inline crypto (VPN, etc.) </td><td> • Storage transport • Transport (TCP / RDMA) • Storage offload • Crypto • Compression • Deduplication • Checksum • RAID </td><td> • Hypervisor I/O virtualization lift and shift • Elastic storage • NVMe™ virtualization </td><td colspan="2"> • Virtual Functions (VFs) / Physical Functions (PFs) / VF switching • Single root I/O virtualization (SR-IOV/ SIOV) • IDE/TDISP • Virtualization adapter • Ethernet/DPOK • NVMe • VirtIO </td></tr><tr><th colspan="5">Security</th></tr><tr><td> • Distributed stateful firewall • Root of Trust </td><td> • DoS Prevention • TLS Proxy </td><td> • Secure Boot • PKE </td><td> • Telemetry • Microsegmentation </td><td> • Confidential Compute </td></tr></table> https://community.amd.com/t5/corporate/amd-pensando-dpu-software/ba-p/630282, CV-AMD0015984.	Software-Defined Networking	Software-Defined Storage		PCIe® Services		• Policy management • Virtual networking • NFV • vSwitch • NAT / load balancer • Inline crypto (VPN, etc.)	• Storage transport • Transport (TCP / RDMA) • Storage offload • Crypto • Compression • Deduplication • Checksum • RAID	• Hypervisor I/O virtualization lift and shift • Elastic storage • NVMe™ virtualization	• Virtual Functions (VFs) / Physical Functions (PFs) / VF switching • Single root I/O virtualization (SR-IOV/ SIOV) • IDE/TDISP • Virtualization adapter • Ethernet/DPOK • NVMe • VirtIO		Security					• Distributed stateful firewall • Root of Trust	• DoS Prevention • TLS Proxy	• Secure Boot • PKE	• Telemetry • Microsegmentation	• Confidential Compute
Software-Defined Networking	Software-Defined Storage		PCIe® Services																		
• Policy management • Virtual networking • NFV • vSwitch • NAT / load balancer • Inline crypto (VPN, etc.)	• Storage transport • Transport (TCP / RDMA) • Storage offload • Crypto • Compression • Deduplication • Checksum • RAID	• Hypervisor I/O virtualization lift and shift • Elastic storage • NVMe™ virtualization	• Virtual Functions (VFs) / Physical Functions (PFs) / VF switching • Single root I/O virtualization (SR-IOV/ SIOV) • IDE/TDISP • Virtualization adapter • Ethernet/DPOK • NVMe • VirtIO																		
Security																					
• Distributed stateful firewall • Root of Trust	• DoS Prevention • TLS Proxy	• Secure Boot • PKE	• Telemetry • Microsegmentation	• Confidential Compute																	

Exhibit E: Claim Chart for U.S. Patent 10,944,634

Claim Language	Infringement Contentions and Exemplary Evidence of Infringement
	 Figure 3. Firewall functions are integrated into the data center fabric, at the ToR. https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/solution-briefs/zero-trust-building-secure-data-centers-cx-10000.pdf, CV-AMD0016035.
[6c] determining, at the tuner server, an adjustment to the transmission	The Accused Products include determining, at the tuner server, an adjustment to the transmission bandwidth is desired. For example, based upon information and belief, the AMD Pensando DPU deployed in the further endpoint device (i.e., tuner server) which is capable of load balancing determines an adjustment to the transmission bandwidth is desired.