

AMD PENSANDO® 2ND GENERATION (“ELBA”) DATA PROCESSING UNIT

April 2024

Overview

The second-generation AMD Pensando™ data processing unit (DPU), “Elba”, offers state-of-the-art capabilities for offloading various data center networking, storage and security services at 2 x 200 Gb/s line rate, at cloud scale. It is source-code compatible with the 1st generation “Capri” DPU, making it easy for existing customers to adopt, and simplifying integration with partner products.

It features a P4-programmable pipeline comprising 144 custom match processing units (MPUs), combined with a 16x A72 Arm® core complex, dedicated data encryption and storage offload engines all tied together via a proprietary fast network-on-a-chip interconnect. Elba’s unique architecture allows AMD, its customers, and partners to create highly efficient and scalable solutions deploying a rich set of software-defined networking, storage, and security features in a virtualized, broadly distributed, centrally managed way.

The Elba DPU is designed to allow the implementation of complex traffic processing and forwarding algorithms, including network virtualization and security, packet encapsulation/decapsulation, inline encryption/decryption, and network address translation.

The Elba DPU architecture supports stateful packet processing, which can be used for flow monitoring, flow-aware security policies implementation, help against DDoS attacks, and more.

Dedicated offload engines facilitate seamless and efficient processing of computation-heavy operations used in various encryption and storage algorithms, including compression/decompression, checksums, RAID, and deduplication hashing.

AMD Pensando DPUs are programmed using the industry-standard P4 language, facilitating the implementation of a wide variety of system solutions. It supports software-defined networking and storage protocols, including NVMe virtualization and transport, and is designed to give developers the agility to develop and deploy new features and modifications throughout the product life cycle.

The Elba DPU form factor and power profile are designed to support multiple system level implementations ranging from a half-height, half-length PCIe card that can fit into the power and cooling profiles of any standard server, to network and security appliances and Smart Switches. The P4-programmable design enables these applications to dynamically re-configure the data processing inside the DPU.

AMD PENSANDO® 2ND GENERATION (“ELBA”) DATA PROCESSING UNIT

Key Features

Feature	Description
Network Interfaces	56 Gb/s PAM4 SerDes supporting Dual 200 GE, Quad-100/50/25/10 GE 1 GE management port
PCIe® Interface	32 lanes of PCIe Gen4, configurable as root complex or end-point mode 2x16 / 4x8 / 8x4 with QoS support in multi-host applications
Data Pipeline	P4 data pipeline comprising 144 match processing units (MPUs) @ 2 GHz Provides high performance capabilities in packet and message processing, at line rate
SOC	16x Arm A72 CPU cores @ 3 GHz with I-Cache, D-Cache, and LLC Cache QSPI Flash, EMMC storage for embedded OS Secure Boot with hardware root of trust
Memory	Dual DDR4-3200 interfaces supporting 8 GB - 64 GB system memory
Offloads	Inline IPsec and DTLS, bulk crypto, PKE, compression, decompression, checksums, deduplication, erasure coding
Scale	2K VNICs, 16 M hardware queues, highly scalable P4 tables (stateful and stateless) accessible at every stage of the pipeline
Scheduling	<ul style="list-style-type: none"> - Queue Group scheduling with Min/Max rate - High Priority Option per queue group - P4 meters and QoS priorities
RoCEv2	Memory-based scatter/gather lists (SGL) in DMA commands Latency optimized hardware data path with all context on DPU
Single Wire Management	Connects 1 GE BMC controller and uplink ports in standby power mode

AMD PENSANDO® 2ND GENERATION (“ELBA”) DATA PROCESSING UNIT

Typical System Configurations

The Elba DPU can be deployed in various system configurations. Typical configurations are listed below.

Model	Power	PCIe	Notes
Standard PCIe Card	25-35 W	8 lanes Gen 3/4	Low Power PCIe card Uplinks: 2x10/25/100 GE + 1x1 GE Mgmt
High Performance PCIe Card	50 W	16 lanes Gen4	Single- or multi-host support Uplinks: 2x200 GE + 1x1 GE Mgmt
Data Center Appliance / Smart Switch	50 W ¹	16 lanes Gen4	Uplinks: 2x200 GE

Key Applications

Application	Details
Advanced Networking	Full support for SDN, virtual private networks (network overlays), L3 ECMP, load balancing, NAT, PAT
Cutting-Edge Security Features	Stateful firewall, security groups, stateless and reflexive ACLs, VPN termination (IPsec), TLS/DTLS encryption, TLS proxy
Enhanced Storage	Full support for SDS, NVMe virtualization, NVMe-oF with RDMA or TCP Transport, AES-XTS data-at-rate encryption, compression/decompression, SHA-3 deduplication, CRC-64/CRC-32 and checksum acceleration
Advanced Observability	Flow-based packet telemetry, stateful connection tracking, latency metrics, highly scalable per-packet counters
Cloud Infrastructure	Virtual private cloud (VPC), security groups, firewalls, DDoS protection, transit gateway and VPC peering, NAT gateway, load balancer, metering, rate control and QoS marking, SR-IOV for workload traffic/resource isolation, IPsec/VPN gateway

¹ DPU only; for full appliance or Smart Switch power figures, refer to the data sheet for the specific product.

AMD PENSANDO® 2ND GENERATION (“ELBA”) DATA PROCESSING UNIT

Hardware and Support

For information about ordering Distributed Services Cards based on the AMD Pensando 2nd generation (“Elba”) DPU, and the appropriate support for your deployment, please contact your AMD sales representative, referencing the Distributed Services Card part number listed below.

Part Number / SKU	Description
DSC2-2Q200-32R32F64P-S4	DSC2-200 (Card) - 2 x 200 Gb/s (QFP56) FHHL

In addition to the PCIe card-based form factor, DPU-only deployments are also possible; consult your AMD sales or partner representative for further information.

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](#)

For Developers

AMD offers a Software-in-Silicon Development Kit (SSDK) to facilitate building customer and partner software solutions leveraging the unique capabilities of the Elba DPU and its existing rich set of software service libraries. Visit amd.com/pensando for details.

AMD PENSANDO® 2ND GENERATION (“ELBA”) DATA PROCESSING UNIT**Disclaimer**

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED ‘AS IS.” AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo, Pensando and combinations thereof are trademarks of Advanced Micro Devices, Inc. Arm® is the registered trademark of Arm Limited in the EU and other countries. PCIe® is a registered trademark of PCI-SIG Corporation. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

© 2023–2024 Advanced Micro Devices, Inc. All Rights Reserved.

amd.com/pensando

PPB23004