

Versal™ Premium Series

Breakthrough Integration of Networked, Power-Optimized Cores on an Adaptable Platform

OVERVIEW

Versal Premium series features breakthrough integration of high-bandwidth, power-optimized networking IP cores with High-Speed Crypto Engines for the fastest, most secure networks.

As a heterogeneous compute platform, the Versal Premium series is engineered to help users reach the highest levels of acceleration for a wide range of compute-intensive data center workloads by providing the highest compute density, custom memory hierarchy, and massive on-chip memory.

With Vivado™ Design Suite and the Vitis™ unified software development platform, the Versal Premium series offers a complete solution stack for hardware and software developers for maximum productivity.

HIGHLIGHTS

Enabling the Fastest, Most Secure Networks


- > 112G PAM4 transceivers for next-gen 800G networks
- > 600G channelized multirate Ethernet cores
- > 600G Interlaken cores with FEC for chip-to-chip interconnect
- > 400G High-Speed Crypto (HSC) Engines for security

Highest Compute Density with Adaptable Acceleration

- > Industry's highest logic density at 7nm for differentiation and adaptability
- > AI/ML for network intelligence, e.g., anomaly detection and self-provisioning
- > Industry's highest DSP compute at 7nm for diverse workload acceleration
- > Integrated PCIe® Gen5 for host CPU-to-accelerator bandwidth
- > Dynamic Function eXchange (DFX) for dynamic workload provisioning

Highly Integrated HW/SW Platform for Greater Productivity

- > Programmable network on chip (NoC) for guaranteed QoS
- > Integrated shell for streamlined connectivity for cloud infrastructure
- > Greatly simplified system design by massive IP integration
- > Vivado Design Suite for hardware developers
- > Vitis development platform for software developers and data scientists



ADAPTABLE TO ANY WORKLOAD

Wired Communications

- > Metro/Core Transport Networks
- > Data Center Interconnect
- > Security Appliances

Data Center Compute

- > Search
- > Recommendation
- > Video Analytics

Test and Measurement

- > Network Tester
- > Mobile Tester
- > PCIe Protocol Analyzer

Aerospace and Defense

- > Radar systems
- > Avionics

FEATURE	DESCRIPTION
Scalar Engines	<ul style="list-style-type: none"> > Complex algorithm processing and decision-making tasks > Dual-core Arm® Cortex®-A72 application processing unit > Dual-core Arm Cortex-R5F real-time processing unit
Platform Management Controller	<ul style="list-style-type: none"> > Boot and configuration and advanced power and thermal management > Security, safety, and reliability enclave > Integrated platform interfaces and high-speed debug
Adaptable Hardware Engines	<ul style="list-style-type: none"> > Re-architected for higher compute capacity and less place and route > High bandwidth, low latency data movement between engines and I/Os > Programmable memory hierarchy for optimal compute efficiency
Intelligent Engines	<ul style="list-style-type: none"> > Enhanced DSP58 Engines for high-precision floating point and low latency > Up to 99TOPs with INT8 and 23TFLOPs with FP32 of DSP compute bandwidth for acceleration
Programmable Network on Chip	<ul style="list-style-type: none"> > High-bandwidth multi-terabit NoC for guaranteed QoS > Programmable framework memory-mapped access to all resources > Easy IP and kernel placement
On-Chip Memory	<ul style="list-style-type: none"> > Up to 1Gb of tightly coupled memory for performance, power, and latency > Up to 123TB/s¹ of on-chip memory bandwidth, 25X vs. GPU²
112G PAM4 Transceivers	<ul style="list-style-type: none"> > Timed with single-lane 100G deployment in 400G infrastructure > Up to 9Tb/s of serial bandwidth in a smaller area with power efficiency
PCIe® Gen5 with DMA and CCIX	<ul style="list-style-type: none"> > Host CPU-to-accelerator communication for next-generation compute applications > Symmetric/asymmetric access to memory with cache coherent interconnect for accelerators
Integrated 600G Ethernet and 100G Multirate Ethernet Cores	<ul style="list-style-type: none"> > Up to 5Tb/s of scalable Ethernet throughput > Multirate: 400/200/100/50/40/25/10G > Multi-standard: FlexE, Flex-O, eCPRI, FCoE, and OTN
Integrated 600G Interlaken Cores with FEC	<ul style="list-style-type: none"> > Scalable chip-to-chip interconnect from 10Gb/s to 600Gb/s > Integrated RS-FEC for power-optimized error correction
400G High-Speed Cryptography Engines	<ul style="list-style-type: none"> > AES-GCM-256/128 engines > Up to 1.6Tb/s of line rate encryption throughput > 400G of MACsec, IPsec, and bulk encryption per engine

TAKE THE NEXT STEP

For more information about the AMD Versal Premium series, visit www.xilinx.com/versal-premium.

1: Memory bandwidth assumes largest Versal Premium device, all available block RAM and UltraRAM at their maximum rates, 72-bit dual-port configuration

2: "Dissecting the Nvidia Turing T4 GPU via Microbenchmarking" - <https://arxiv.org/pdf/1903.07486.pdf>

DISCLAIMERS

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

COPYRIGHT NOTICE

© Copyright 2023 Advanced Micro Devices, Inc. All rights reserved. Xilinx, the Xilinx logo, AMD, the AMD Arrow logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. AMBA, AMBA Designer, ARM, ARM1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of ARM in the EU and other countries. PCIe, and PCI Express are trademarks of PCI-SIG and used under license. PID# 231846771-I

