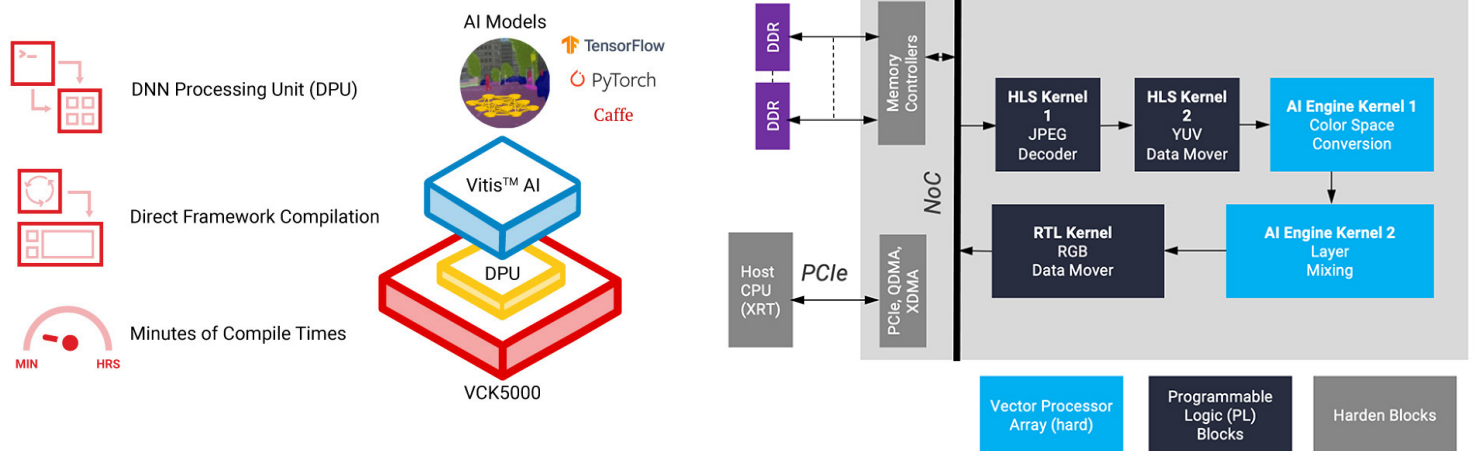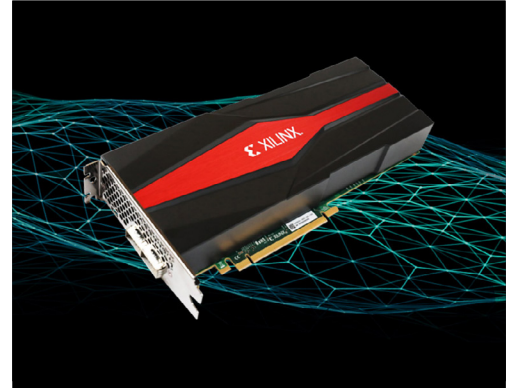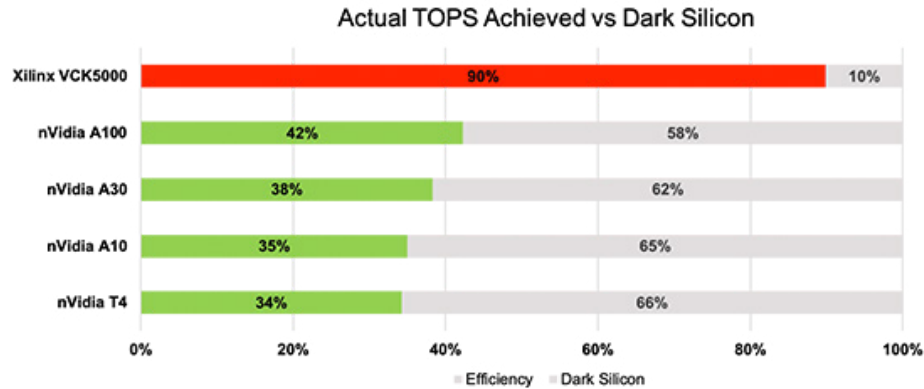# VCK5000
Versal Development Card

OVERVIEW

The Xilinx® VCK5000 Versal® development card is built on the Xilinx 7nm Versal ACAP architecture and is designed for 5G, DC compute, AI, Signal Processing, Radar, and others. The Domain-Specific Architecture brings strong horsepower per watt while keeping ease-of-use in mind with C/C++, software programmability.

Delivering the world's first zero dark silicon (near 100% compute efficiency per watt) in standard AI benchmarks and 2x TCO compared to the flagship nVidia GPUs, the VCK5000 is an ideal development platform for CNN, RNN, and NLP acceleration for your cloud and edge applications.

VCK5000 is fully supported by Vitis™ unified software platform, Vitis AI™ development environments, and partner AI solutions such as Mipsology and Aupera. If you are an AI developer, bring your TensorFlow / Pytorch trained models directly to infer using Vitis AI or Mipsology Zebra. If you are looking to implement algorithm acceleration with AI engine and programmable logic, we provide AI engine C/C++ high-level abstraction APIs and Vitis Accelerated Libraries. The Vitis end-to-end flow is developed using C/C++ to run on X86 or embedded processors and manage runtime interactions with the accelerator with XRT. The hardware component, or kernel, can be developed using C/C++, or RTL target on PL or AI Engines.

The World's First "0 Dark Silicon" AI Accelerator

Actual TOPS Achieved vs Dark Silicon

	Efficiency	Dark Silicon
Xilinx VCK5000	90%	10%
nVidia A100	42%	58%
nVidia A30	38%	62%
nVidia A10	35%	65%
nVidia T4	34%	66%

Source: https://developer.Nvidia.com/deep-learning-performance-training-inference

HIGHLIGHTS

AI Inference Development

2x TCO vs Mainstream nVidia GPUs

> 2x perf/w and perf/$ compared to nVidia GPUs with standard MLPerf models

> Archives near 100% compute efficiency

> Consume less than 100W at card level

2x End-to-End Video Analytics Throughput vs nVidia GPUs

> Full pipeline from H.264 decode to computer vision to up to 10 AI models

> Video decode and CV run on x86 CPU or discrete U30 Alveo card

> Plug-in based pipeline composition from FFmpeg / Gstreamer

Easy to Use with Familiar Frameworks

> Easy-to-use software flow for any CPU & GPU users, no hardware programming required

> Run inference from Tensorflow framework directly on board

> State-of-the-art model supported with mainstream frameworks Pytorch, Tensorflow, Tensoflow 2 and Caffe

AI Engine Development

Power and Performance

> Up to 10x performance improvements compared to previous generation Xilinx UltraScale+™ with less power in diverse applicationsl

Software Familiarity

> A familiar software development flow with Vitis unified software platform

> Accelerate your applications faster with AI engine C/C++

Mixed Kernel Development

> Customize your own data pipeline with mixed kernels

> Develop AIE kernels in C/C++, PL kernels in RTL or HLS, and let Vitis stitch together the full system

AMD
XILINX

SPECIFICATION

CARD SPECIFICATIONS	VCK5000	
Device	VC1902	
Compute	Active	Passive*
Peak INT8 TOPS	145	145
DDR Memory Bandwidth	70	70
Internal SRAM Bandwidth	37	37
Look-up Tables (LUTS)	899,840	899,840
PCI Express	12	12
Dimensions		
Height	Full	Full
Length	Full	3/4
Width	Dual Slot	Dual Slot
Memory		
Off-chip Memory Capacity	16 GB	16 GB
Off-chip Total Bandwidth	102.4 GB/s	102.4 GB/s
Internal SRAM Capacity	23.9 MB	23.9 MB
Internal SRAM Total Bandwidth	23.5 TB/s	23.5 TB/s
Interfaces		
PCI Express	Gen3 x 16 / Gen4x 8	Gen3 x 16 / Gen4x 8
Network Interfaces	2x QSFP28 (100GbE)	2x QSFP28 (100GbE)

TAKE THE NEXT STEP

Learn more about Xilinx VCK5000   www.xilinx.com/products/boards-and-kits/vck5000.html

Purchase the VCK5000 Development Card for $2,745

Get started with Xilinx partner solutions on VCK5000

Get started with Xilinx Vitis AI solution