

VPX3U-AD5000E-CX7-OPTICAL

HPEC with NVIDIA Ada RTX™ 5000, ConnectX-7 SmartNIC & 40/100GbE Optical

OVERVIEW

The VPX3U-AD5000E-CX7-OPTICAL HPEC module includes an NVIDIA RTX™ 5000E (AD103) embedded GPU and a ConnectX®-7 SmartNIC. The NVIDIA RTX 5000E embedded GPU provides advanced processing capabilities for high performance embedded computing (HPEC) and artificial intelligence (AI) processing. The ConnectX-7 provides the PCIe connectivity and dual 40/100GBASE optical fabric interfaces (P2A and P2B) needed to move large datasets efficiently.

The NVIDIA Ada architecture includes CUDA cores for HPEC, 4th gen Tensor cores for AI computations, and 3rd gen (RT) cores for visually accurate rendering. The Ada GPU uses a TSMC 4N NVIDIA Custom Manufacturing Process with increased efficiency. The denser Ada GPUs have more CUDA and Tensor cores operating at higher clock frequencies, delivering significantly more performance per watt compared to the previous generation product. For example, the 153W can achieve two to three times the performance per watt of the 144L (depending on power and temperature).

The NVIDIA ConnectX-7 SmartNIC provides PCIe connectivity and 40/100GBASE optical fabric interfaces. ConnectX-7 is ideal for the high-speed, secure, data transfer capabilities required for data-heavy tasks such as sensor data processing and other C5ISR tasks. The ConnectX-7 also provides support for RDMA over Converged Ethernet (RoCE), enabling an efficient method for transferring data across the network to the GPU, including support for NVIDIA GPUDirect RDMA.

KEY FEATURES

- NVIDIA RTX™ 5000E (AD103) GPU with 9728 CUDA Cores, 304 Tensor Cores, 76 RT Cores
- NVIDIA® ConnectX®-7 provides the module with up to 100GbE Ethernet and a configurable PCIe switch, with 40/100GBASE optical fabric on P2A & P2B
- 16 GB GDDR6 256-bit VRAM with ECC support
- Module power: 100W - 140W, configurable

GPU FEATURES

- Ada GPGPU parallel processing:
 - CUDA Toolkit 12, CUDA Compute capability 8.9
 - CUDA-X AI and CUDA-X HPEC libraries
 - OpenCL™ 3.0, DirectX® 12 Ultimate, OpenGL 4.6, OpenGL ES 3.2, Vulkan™ 1.2
- 304 Tensor Cores (4th Gen)
- 76 Ray Tracing cores (3rd Gen)
- NVENC (8th Gen) and NVDEC (5th Gen) with up to 8K video encoding and hardware decoding support



WOLF

nvidia

CONNECTIVITY / SYSTEM MANAGEMENT

- On-board IPMI controller for system management
- PCIe configurable switch
- Switching is offloaded from the CPU to the ConnectX with NVIDIA ASAP2 technology
- Daisy Chain options
- Support for dual 40/100GBASE optical fabric links (P2A & P2B)
- GPUDirect RDMA and RoCE support
- Block-level hardware encryption and the use of dedicated encryption keys per user
- Linux and Windows drivers

MECHANICAL/OPEN SYSTEMS ARCHITECTURE

- High level of ruggedization:
 - Rugged conduction cooled
 - Operating temperature: -40°C to +70°C standard, operational to +85°C
 - Vibration (sine wave): 10G peak, 5 - 2000Hz
 - Shock: 40G peak
- Dimensions: 160mm x 100mm x 25.4mm
- Weight (approximately): 1.3 kg
- SOSA® Aligned Payload 14.6.11, P2A & P2B Optical


THIS INFORMATION IS SUBJECT TO CHANGE

VPX3U-AD5000E-CX7-OPTICAL

HPEC with NVIDIA Ada RTX™ 5000, ConnectX-7 SmartNIC & 40/100GbE Optical

CHIP-DOWN DESIGN

The VPX3U-AD5000E-CX7-OPTICAL module uses a WOLF chip-down design to provide NVIDIA's advanced Ada architecture GPU and ConnectX-7 SmartNIC technology on an extremely rugged board, making it an excellent choice for aerospace and defense applications. WOLF designs and manufactures modules in North America with full component traceability.



POWER AND PERFORMANCE

An NVIDIA GPU will operate at the GPU clock speed available at the set TGP (total GPU subsystem power). The highest clock speeds are available at the highest TGP power allowed by the GPU. When the TGP setting is decreased the clock speed will also decrease resulting in a decrease in processing speed. The GPU clock speed will also decrease if the GPU temperature exceeds 87°C to protect the GPU from heat damage. If the GPU temperature is below 86.5°C the GPU will operate at boost clock speeds when higher processing is required.

The Ada AD103 GPU in this 3U VPX module will default to a TGP power of 80W. At a TGP of 80W the base clock of 930 MHz provides up to 18.1 TFLOPS, and at higher GPU loads the boost clock can provide up to 1680 MHz which provides up to 32.7 TFLOPS. A higher TGP can be configured if the GPU can be cooled sufficiently, while a lower TGP can be configured when operating in hotter environments.

THIS INFORMATION IS SUBJECT TO CHANGE

VPX3U-AD5000E-CX7-OPTICAL

HPEC with NVIDIA Ada RTX™ 5000, ConnectX-7 SmartNIC & 40/100GbE Optical

NVIDIA ADA GPU

NVIDIA Ada GPUs are manufactured using a new TSMC 4N NVIDIA Custom Manufacturing Process, an enhanced version of the N5 (5nm) node process. This allows a higher transistor density and lower voltage requirements, which provides increased efficiency. As a result, Ada GPUs have many more CUDA cores at the same die size as the previous generation, and higher clock speeds at the same power level, which leads to greatly increased processing/watt compared to the previous generation. The new Ada architecture also provides a big increase in the GPU's memory cache size, providing a boost to memory subsystem handling at the same bandwidth.

TENSOR CORES FOR ARTIFICIAL INTELLIGENCE AND HPEC

Tensor Cores are designed to speed up the tensor / matrix computations used for deep learning neural network training and inferencing operations. NVIDIA Ada architecture GPUs include the fourth-generation Tensor Core design which supports many data types for improved performance, efficiency, and programming flexibility, including a sparsity feature, a Tensor Float 32 (TF32) precision mode, and a new FP8 precision mode. NVIDIA provides CUDA-X AI and CUDA-X HPEC libraries which have been designed to work with NVIDIA Tensor Core GPUs to provide the tools needed to accelerate development of applications for AI and HPEC.

CONNECTX-7 WITH PCIE GEN4 AND OPTICAL ETHERNET 40/100GBE

Getting large amounts of data into and out of a module is an important system design consideration. The WOLF-153W module includes a ConnectX-7 SmartNIC, which provides a configurable PCIe interface. It also provides dual 40/100GBASE optical fabric interfaces (P2A and P2B), RDMA over Converged Ethernet (RoCE) with support for NVIDIA GPUDirect RDMA, and enhanced security features such as hardware-verified secure boot, hardware-accelerated cryptography, and encrypted storage.

OPTICAL FABRIC CONNECTOR

The WOLF-153W implements its dual optical fabric interfaces (P2A and P2B) using the Smiths Interconnect LightCONEX® 28G LC Series active blind-mate optical plug-in module connector (Style C). This fabric interface supports 28 Gbps per channel, is offered in a 4-lane (4TRX) 100G full-duplex configuration, and supports operation from -40°C to +85°C. This connector capability maps directly to the 40/100GBASE optical links shown on P2A and P2B in the 153W, enabling high-bandwidth optical networking via the on-board ConnectX-7.



Figure 1: LightCONEX 28G 4TRX - Style C

SOSA SLOT PROFILE SUPPORT

This module's configurable switch provides support for a SOSA aligned payload slot profile intended for optical fabric connectivity. The following SOSA aligned profile is supported:

- 14.6.11-0 Payload Slot Profile, P2A & P2B Optical

THIS INFORMATION IS SUBJECT TO CHANGE

VPX3U-AD5000E-CX7-OPTICAL

HPEC with NVIDIA Ada RTX™ 5000, ConnectX-7 SmartNIC & 40/100GbE Optical

ORDERING CODES

The following table defines series of common order codes for the VPX3U-AD5000E-CX7-OPTICAL module. The asterisks denote characters of the part number that are defined based on common configuration options. Some configuration options for this module include:

- Default Power Threshold
- PCIe Configuration Options
- Variant Locked
- Conformal Coating
- Optical Ethernet (P2A/P2B)

ORDERING NUMBER	DESCRIPTION
3U VPX Ada RTX5000E with ConnectX-7 SmartNIC and 40/100GBE Optical Single Slot Configurations	
153W33-F***-***VPX3vA0	3U VPX, Conduction Cooled, NVIDIA Ada RTX5000, ConnectX-7, SOSA Payload profile 14.6.11 with P2A & P2B Optical
Related Product:	
153L33-F***-***VPX3vA0	3U VPX, Conduction Cooled, NVIDIA Ada RTX5000, ConnectX-7, SOSA Payload profile with P2 depopulated
153L33-F***-***VPX3vA0	3U VPX, Conduction Cooled, NVIDIA Ada RTX5000, ConnectX-7, SOSA Payload profile with P2B depopulated

* Contact WOLF to determine the appropriate configuration for your system.

MANUFACTURING & QUALITY ASSURANCE

WOLF designs modules to pass the following environmental standards:

- MIL-STD-810 (United States Military Standard for Environmental Engineering Considerations and Laboratory Tests)
- MIL-HDBK-217 (Reliability Prediction of Electronic Equipment)
- RTCA DO-160 (Environmental Conditions and Test Procedures for Airborne Equipment) on request

WOLF complies with the following quality management systems:

- AS9100D: Quality Management System - Requirements for Aviation, Space and Defense Organizations (certified)
- ISO 9001:2015: Quality management systems (certified)
- AS5553: Counterfeit Electronic Parts; Avoidance, Detection, Mitigation, and Disposition (compliant)
- NIST SP 800-171: Protecting Controlled Unclassified Information in Nonfederal Systems (compliant)

Boards are manufactured to meet the following standards:

- IPC-A-610 CLASS 3 (Acceptability of Electronic Assemblies)
- IPC 6012 CLASS 3 (Qualification and Performance Specification for Rigid Printed Boards, Class 3 for High Reliability Electronic Products)
- IPC J-STD-001 (Requirements for Soldered Electrical and Electronic Assemblies)



THIS INFORMATION IS SUBJECT TO CHANGE