

HPC with NVIDIA Blackwell RTX 5000 GPU and ConnectX-7

KEY FEATURES

- NVIDIA RTX™ 5000 (GB203) GPU with 10496 CUDA Cores, 320 Tensor Cores
- NVIDIA® ConnectX®-7 provides the module with up to 100GbE Ethernet and a configurable PCIe switch
- 24 GB GDDR7 256-bit VRAM with ECC support
- Bandwidth: Up to 896 Gbps (50% more than ADA)
- Module power: 100W - 180W, configurable

GPU FEATURES

- Blackwell GPGPU parallel processing:
 - CUDA Toolkit 12, CUDA Compute capability 10.0
 - OpenCL™ 3.0, DirectX® 12 Ultimate, OpenGL 4.6, OpenGL ES 3.2, Vulkan™ 1.2
- 5th Gen Tensor Cores with new data precisions (new: FP4 and FP6, FP8 Gen2)
- NVENC (9th Gen) and NVDEC (6th Gen) with up to 8K video encoding and hardware decoding support

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
- Support for 40/100GBASE-KR4 protocols
- 10/25GBASE-KR Data and Control planes
- 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 Random: VITA 47.1 Class V3 (5 to 2000Hz)
 - Vibration Sine: 10g peak (5 to 2000 Hz)
 - Shock: 40G (MIL-STD-810H, Method 516.8)
- Dimensions: 160mm x 100mm x 25.4mm
- Weight (approximately): 1.6 kg
- SOSA™ Aligned slot profile: 14.6.11-0 or 14.6.13-0

OVERVIEW

The VPX3U-BW5000E-CX7 HPEC module includes an NVIDIA RTX 5000 Blackwell embedded GPU and a ConnectX SmartNIC. The NVIDIA RTX 5000 Blackwell embedded GPU provides the advanced processing capabilities for high performance embedded computing (HPC) and artificial intelligence (AI) processing. The ConnectX-7 provides the Ethernet and PCIe connectivity needed to move large datasets efficiently.

The NVIDIA Blackwell architecture includes CUDA cores for HPC, and 5th generation Tensor cores for AI and data science computations. The Blackwell GPU has an improved architecture which provides increased efficiency. The module also supports 24GB of GDDR7 memory which provides over 50% higher bandwidth compared to the previous generation. The GPU supports PCIe up to x16 providing a fast data transfer path to/from the module.

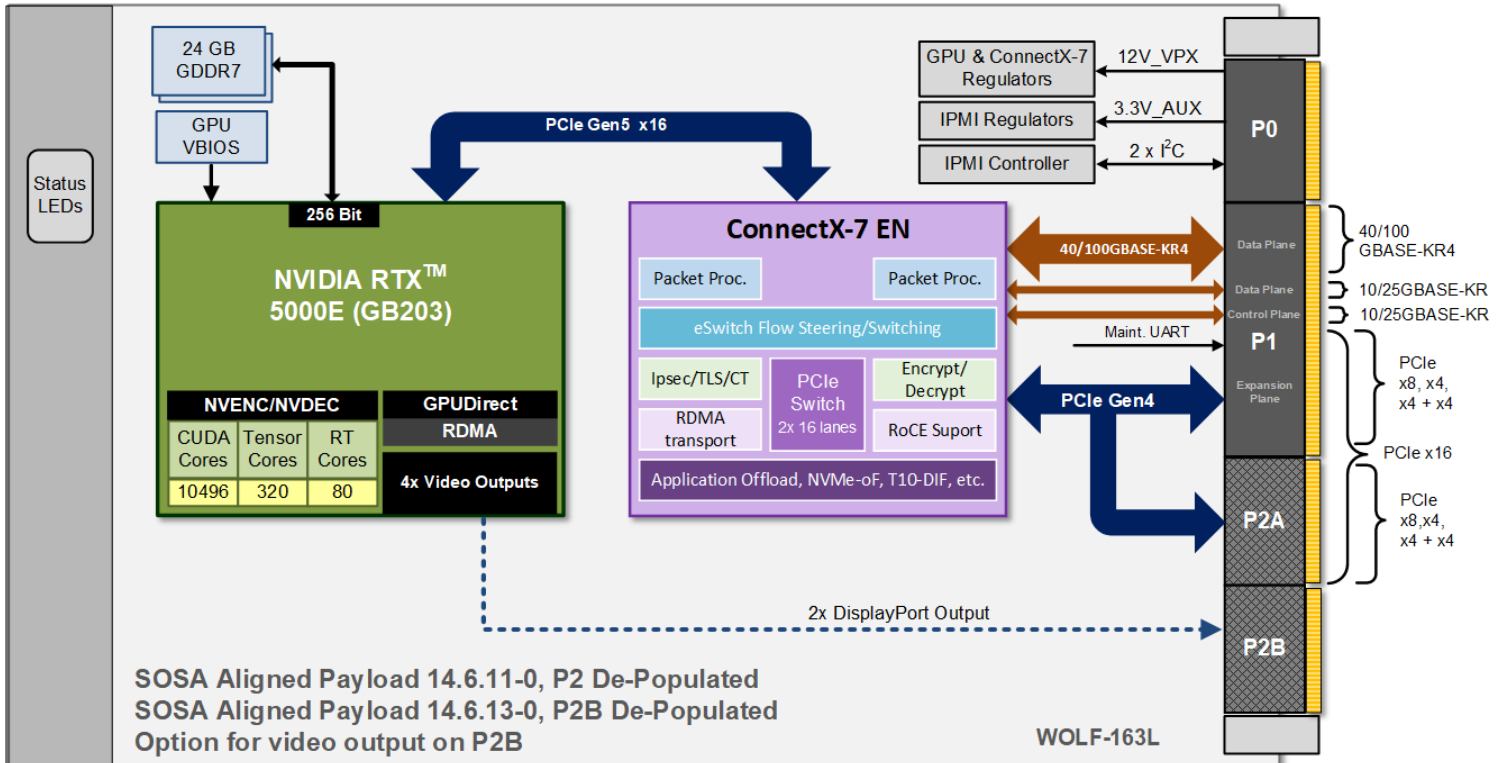
The NVIDIA ConnectX-7 SmartNIC provides PCIe and Ethernet connectivity. ConnectX-7 is ideal for the high-speed, secure, data transfer capabilities required for data-heavy tasks such as sensor data processing and other C5ISR tasks. The ConnectX-7 also provides support for RDMA over Converged Ethernet (RoCE), enabling the fastest method for transferring data across the network to the GPU.



This information is subject to change

CHIP-DOWN DESIGN

The VPX3U-BW5000E-CX7 module uses a WOLF chip-down design to provide NVIDIA’s advanced Blackwell architecture GPU and ConnectX 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 clock speed is dependent on the TGP (total GPU subsystem power) and the GPU temperature. 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. If the GPU temperature exceeds 87°C the GPU clock speed will also decrease to protect the GPU from heat damage. If the GPU temperature is below 86.5°C the GPU can operate at maximum boost clock speeds at the currently available power when the GPU detects that higher processing is required.

The Blackwell GB203 GPU in this 3U VPX module will default to a TGP power of 100W. At a TGP of 100W the base clock of 1125 MHz provides up to 24.2 TFLOPS, and at higher GPU loads the boost clock can run up to 1792 MHz which provides up to 38.5 TFLOPS. A higher TGP of up to 150W can be configured if the GPU can be cooled sufficiently, with a maximum boost clock of 2370 MHz providing up to 50 TFLOPS.

NVIDIA BLACKWELL GPU

NVIDIA Blackwell GPUs have an improved architecture which provides increased efficiency. Blackwell GPUs have CUDA cores that can all handle either FL or INT operations, whereas previous generations restricted half of the CUDA cores for FL operations only. This is important for tasks that require lots of small, frequent lookups resulting in frequent address computations, which are commonly required by the matrix operations used by AI and HPC calculations. Optimizing these lookups ensures the data can reach the processing cores efficiently, permitting full use of the processing power. The new Blackwell architecture also supports GDDR7 memory, which provides 55% more memory bandwidth. With the increased performance and memory handling abilities, and improved next Gen Tensor cores, the Blackwell GPUs are able provide significant performance increases compared to the previous generation.

TENSOR CORES FOR ARTIFICIAL INTELLIGENCE AND HPC

Tensor Cores are designed to speed up the tensor / matrix computations used for deep learning neural network training and inferencing operations. NVIDIA Blackwell architecture GPUs include the fifth-generation Tensor Core design which supports many data types for improved performance, efficiency, and programming flexibility, including support for new INT4 and INT6 precision modes and microscaling formats. NVIDIA provides CUDA-X AI and CUDA-X HPEC libraires designed to work with NVIDIA GPUs to provide the tools needed to accelerate development of applications for AI and HPEC.

CONNECTX-7 WITH PCIe AND ETHERNET 100GbE SUPPORT

Getting large amounts of data into and out of a module is an important system design consideration. The WOLF-163L module includes a ConnectX-7 SmartNIC, which provides a configurable PCIe interface. It also provides up to 100GBASE-KR4 on the data plane, 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.

HARDWARE ACCELERATED VIDEO ENCODE / DECODE

The Blackwell GPU includes the NVENC video encode and NVENC decode hardware acceleration engine. Using the GPU for video encoding provides an efficient, high quality method to achieve real time 8K and 4K encoding without burdening the system CPU. The Blackwell encoding engine includes support for several popular codecs and is the first GPU to include AV1 hardware encoding and decoding support. The NVIDIA Video Codec SDK provides a complete set of APIs, samples and documentation for hardware accelerated video encode and decode.

SOSA SLOT PROFILE SUPPORT

This module's configurable switch provides support for SOSA aligned slot profiles. The module can be configured with pin mappings that are compatible with previous generation WOLF-153L and 144L modules, which allows the WOLF-163L to be a plug-in upgrade for previous WOLF products.

The following SOSA aligned profiles are supported:

- 14.6.11-0 Payload Slot Profile, P2 depopulated
- 14.6.13-0 Payload Slot Profile, P2B depopulated
- Option for video output on P2B

This information is subject to change

ORDERING CODES

The following table defines series of common order codes for the VPX3U-BW5000E-CX7 module. The asterisks denote characters of the part number that are defined based on common configuration options. Some configuration options for this module include:

- Payload slot profile
- PCIe Configuration Options
- Conformal Coating
- Default Power Threshold
- Ethernet configuration
- Variant Locked

Ordering Number	Description
3U VPX Blackwell RTX5000 with ConnectX-7 SmartNIC Configurations	
163L33-F***-***VPX3vA0	3U VPX, Conduction Cooled, NVIDIA Blackwell RTX5000, ConnectX-7, SOSA Payload profile 14.6.11-0 with P2 depopulated, PCIe Gen4 x8, 3 network ports: 100GbE and two 25 GbE
163L33-F***-***VPX3vA0	3U VPX, Conduction Cooled, NVIDIA Blackwell RTX5000, ConnectX-7, SOSA Payload profile 14.6.13-0 with P2B depopulated, PCIe Gen4 x16, 3 network ports: 100GbE and two 25 GbE

* Contact Sales for the latest Ordering Numbers and available options.

MANUFACTURING AND 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 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)



Datasheet Rev.9

WOLF-163L

This information is subject to change

www.WOLFAdvancedTechnology.com

sales@wolf-at.com | International: +1 (905) 852-1163 | USA Toll-Free: 1 (800) 931-4114