

HPEC with NVIDIA Ampere RTX A4500 and ConnectX-7

KEY FEATURES

- NVIDIA RTX™ A4500 (GA104) GPU with 5888 CUDA Cores, 184 Tensor Cores, 46 RT Cores
- NVIDIA® ConnectX®-7 provides the module with up to 100GbE Ethernet and a PCIe Gen4 switch
- 16 GB GDDR6 256-bit memory with up to 512 GB/s
- Module power: 80-170W, configurable

GPU FEATURES

- Ampere GPGPU parallel processing:
 - CUDA Toolkit 11, CUDA Compute capability 8.6
 - CUDA-X AI and CUDA-X HPEC libraires
 - OpenCL™ 3.0, DirectX® 12 Ultimate, OpenGL 4.6, OpenGL ES 3.2, Vulkan™ 1.2
- 184 Tensor Cores (3rd Gen)
- 46 Ray Tracing cores (2nd Gen)
- NVENC (7th Gen) and NVDEC (5th Gen) with up to 8K video encoding and hardware decoding support

CONNECTIVITY / SYSTEM MANAGEMENT

- On-board IPMI controller for system management
- PCIe Gen5 configurable switch
- Switching is offloaded from the CPU to the ConnectX with NVIDIA ASAP² technology
- Daisy Chain options
- Backplane Ethernet with 10/40/100 GBASE-KR4 protocol support
- GPUDirect RDMA support
- Block-level hardware encryption and the use of dedicated encryption keys per user
- Windows and Linux drivers

MECHANICAL / OPEN SYSTEMS ARCHITECTURE

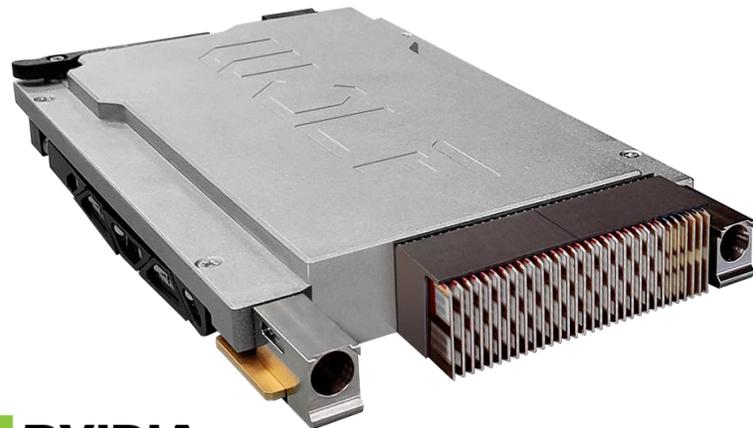
- High level of ruggedization:
 - Rugged conduction cooled
 - Operating temperature: -40° to +85°C
 - Vibration (sine wave): 10G peak, 5 - 2000Hz
 - Shock: 40G peak
- Dimensions: 160mm x 100mm x 25.4mm
- Weight (approximately): 1.5 kg
- SOSA Aligned options with depopulated P2 or P2B
- Support for SOSA 14.6.11 or 14.6.13 slot profile

OVERVIEW

The VPX3U-A4500E-CX7 HPEC module includes an NVIDIA RTX™ A4500 embedded GPU and a ConnectX SmartNIC. The NVIDIA RTX™ A4500 embedded GPU provides the advanced processing capabilities for high performance embedded computing (HPEC) and artificial intelligence (AI) processing. The ConnectX-7 provides the Ethernet and PCIe connectivity needed to move large datasets efficiently.

The NVIDIA Ampere architecture includes CUDA cores for HPEC, 3rd generation Tensor cores for AI and data science computations, and 2nd generation Ray Tracing (RT) cores for visually accurate rendering. The Ampere architecture introduces support for PCIe Gen 4, improvements to memory handling, and improvements to warp scheduling, which all optimize the way data is distributed for processing. The Ampere GPU is very power efficient, delivering significantly more performance per watt compared to previous generations.

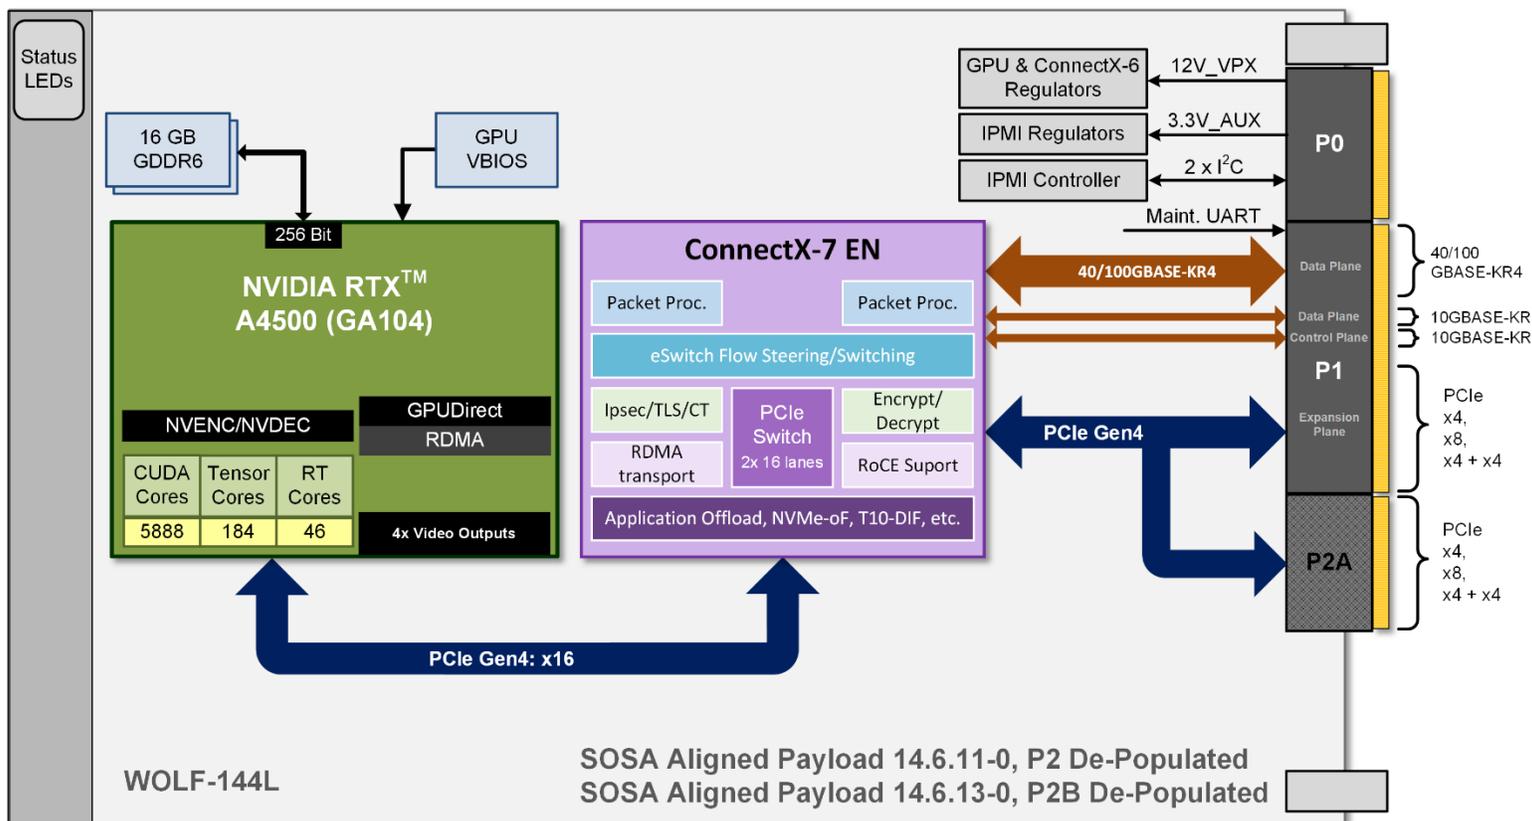
The NVIDIA ConnectX-7 SmartNIC provides PCIe and Ethernet connectivity. ConnectX-7 is ideal for the high-speed, secure, data transfer capabilities required for data-heavy tasks such as sensor data processing and other C5ISR tasks. The ConnectX-7 also provides support for RDMA over Converged Ethernet (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-A4500E-CX7 module uses a WOLF chip-down design to provide NVIDIA’s advanced Ampere 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 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 base clock speed will also decrease if the GPU temperature exceeds 89°C to protect the GPU from heat damage. If the GPU temperature is below 87°C the GPU can operate at full boost clock speeds when higher processing is required.

The Ampere GA104 GPU typically operates at TGP power levels from 80W to 115W, although it can be configured to a lower maximum power. At 115W the base clock of 930 MHz provides 11 TFLOPS FP32 performance while the boost clock of 1500 MHz provides 17.7 TFLOPS. At 80W the base clock of 510 MHz provides 6.0 TFLOPS while the boost clock of 1215 MHz provides 14.3 TFLOPS.

This information is subject to change

NVIDIA AMPERE STREAMING MULTIPROCESSOR (SM)

Each NVIDIA Ampere architecture streaming multiprocessor (SM) partition contains CUDA cores for FP and INT operations, Tensor cores for AI, Ray Tracing (RT) cores for rendering, Texture Units, a register file, and L1/Shared Memory. Each previous generation Turing SM partition had two primary datapaths, with one able to process FP32 operations while the other was limited to integer operations. An Ampere SM partition's two primary datapaths can both process FP32 operations, with one datapath dedicated to FP32 operations and the other capable of executing either FP32 or integer operations. For operations which require only FP32 operations this doubles the number of available CUDA cores per SM. This change to the available functionality of the primary datapaths along with many other improvements to the other components in the Streaming Multiprocessor allows Ampere GPUs to provide significant performance improvements.

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 Ampere architecture GPUs include the third-generation Tensor Core design which supports many new data types for improved performance, efficiency, and programming flexibility, including a new sparsity feature and a new Tensor Float 32 (TF32) precision mode.

NVIDIA provides CUDA-X AI and CUDA-X HPEC libraires 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 ETHERNET 100GbE AND PCIe GEN4

Getting large amounts of data into and out of a module is an important system design consideration. The WOLF-144L module includes a ConnectX-7 SmartNIC, which provides a configurable PCIe Gen4 switch. 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 Ampere GPU includes the NVENC video encode (version 7.2) and NVENC decode (version 5) hardware acceleration engine. Using the Ampere GPU for video encoding provides an efficient, high quality method to achieve real time 8K and 4K encoding without burdening the system CPU. The Ampere decoding engine includes support for several popular codecs and is the first GPU to include AV1 hardware 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 older generation WOLF-134C and WOLF-134S modules, which allows the WOLF-144L to be a plug-in upgrade for previous WOLF products.

The following SOSA aligned profiles are supported:

- SLT3-PAY-1F1U1S1S1U1U2F1H-14.6.11-0 Payload Slot Profile, P2 depopulated
- SLT3-PAY-1F1U1S1S1U1U4F1J-14.6.13-0 Payload Slot Profile, P2B depopulated

ORDERING CODES

The following table defines series of common order codes for the VPX3U-A4500E-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:

- Default Power Threshold
- PCIe Configuration Options
- Variant Locked
- Conformal Coatings
- Ethernet configuration
- P2 or P2B de-populated

| Ordering Number | Description |
|---|---|
| 3U VPX Ampere A4500 with CX7 Switch Single Slot Configurations | |
| 144L33-F**0-***VPX3vA0 | 3U VPX, Conduction Cooled, NVIDIA Ampere A4500, ConnectX-7, SOSA Payload profile 14.6.11 with P2 depopulated |
| 144L33-F**1-***VPX3vA0 | 3U VPX, Conduction Cooled, NVIDIA Ampere A4500, ConnectX-7, SOSA Payload profile 14.6.13 with P2B depopulated, |
| Related Product: | |
| 144833-F005-***VPX3vA0 | 3U VPX, Conduction Cooled, 1", NVIDIA Ampere A4500, PCIe switch, SOSA Payload profile 14.6.11 with P2 depopulated |

* 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)



Intertek



ITAR Free



Datasheet Rev.11

WOLF-144L