

VPX3U-A4500E-VO

NVIDIA Ampere 17.7 TFLOPS, 160 Tensor Cores, 4 DP Out

KEY FEATURES

- NVIDIA RTX™ GA104 GPU with 17.66 TFLOPS FP32 peak performance*, 5888 CUDA Cores, 184 Tensor Cores, 46 RT Cores, 4 DP outputs
 - 16 GB GDDR6 256-bit memory with up to 512 GB/s
 - PCIe x16 Gen4 with a configurable PCIe switch
 - Module power: TBD-180W, configurable
- *Peak performance requires the highest power configuration mode.

GPU FEATURES

- Four simultaneous DisplayPort 1.4 outputs:
 - 4K at 120Hz or 8K at 60Hz with 10-bit color depth
 - Support for High Dynamic Range (HDR) video
 - HDMI and DVI options
- Ampere GPGPU parallel processing:
 - CUDA Toolkit 11, CUDA Compute capability 8.6
 - OpenCL™ 3.0, DirectX® 12 Ultimate, OpenGL 4.6, OpenGL ES 3.2, Vulkan™ 1.2
- 184 Tensor Cores (3rd Gen), 68 (dense) / 136 (sparse) Tensor TFLOPS
- 46 Ray Tracing cores (2nd Gen)
- Error Correcting Code (ECC) memory
- NVENC (7th Gen) and NVDEC (5th Gen) with up to 8K video encoding and hardware decoding support

CONNECTIVITY / SYSTEM MANAGEMENT

- GPU with PCIe x16 Gen4 interface
- Configurable PCIe Gen4 switch
- Daisy Chain option supported
- Windows and Linux drivers
- On-board IPMI controller for system management

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): TBD
- ANSI/VITA 48, 65 (VPX-REDI, OpenVPX)
- SOSA Aligned options with depopulated P2 or support for the SOSA Legacy payload slot profile


IN DEVELOPMENT

OVERVIEW

The VPX3U-A4500E-VO module includes an NVIDIA RTX™ A4500 embedded GPU in a rugged 3U VPX module. Built on the NVIDIA Ampere architecture this GPU includes CUDA cores for parallel processing, Tensor cores for dedicated AI-accelerated compute, and Ray Tracing cores for superior rendering speeds. This module also includes a configurable PCIe Gen4 switch.

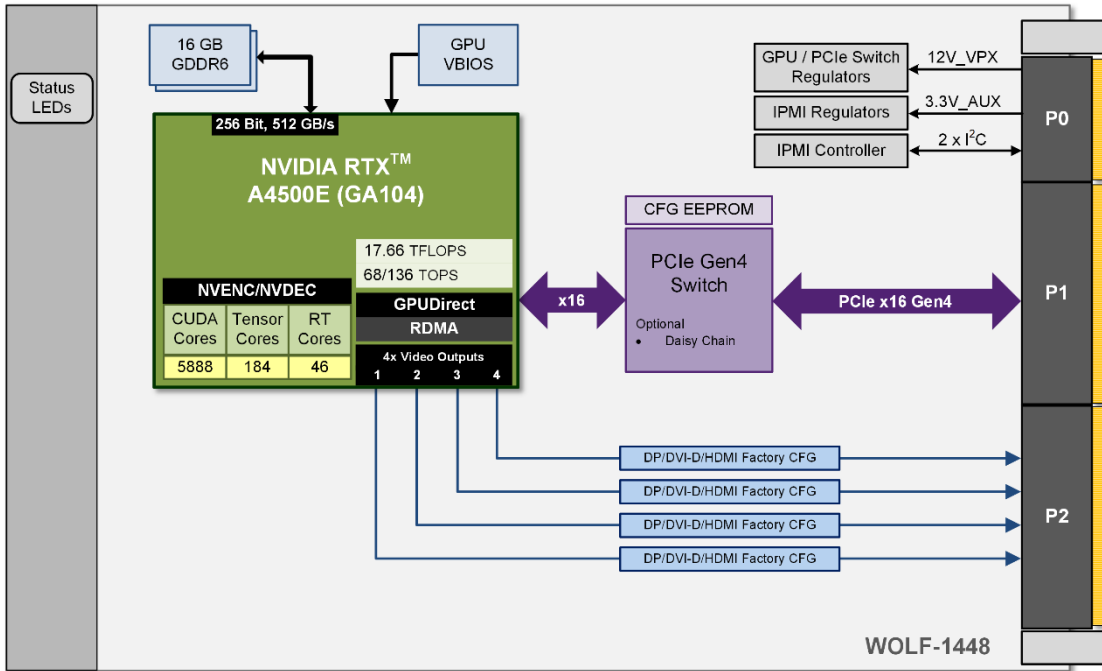
The NVIDIA Ampere architecture has introduced many significant improvements to the performance and efficiency of the GPU, with more flexible CUDA FP32/INT core use, more efficient third generation Tensor cores, and second generation RT cores. The Ampere GPU fabrication uses an 8nm manufacturing processing providing significant power improvements which, along with other Ampere architecture improvements, can provide up to 154 GFLOPS/W, providing almost twice the performance per slot compared to the previous Turing generation's 86 GFLOPS/W or the Pascal generation's 62 GFLOPS/W.

Unlocking the best performance requires the best cooling capability. This module takes advantage of WOLF Advanced Metal, a material which provides lower thermal resistance between the GPU and the wedgelocks, allowing the GPU to run at higher power levels to achieve higher performance.



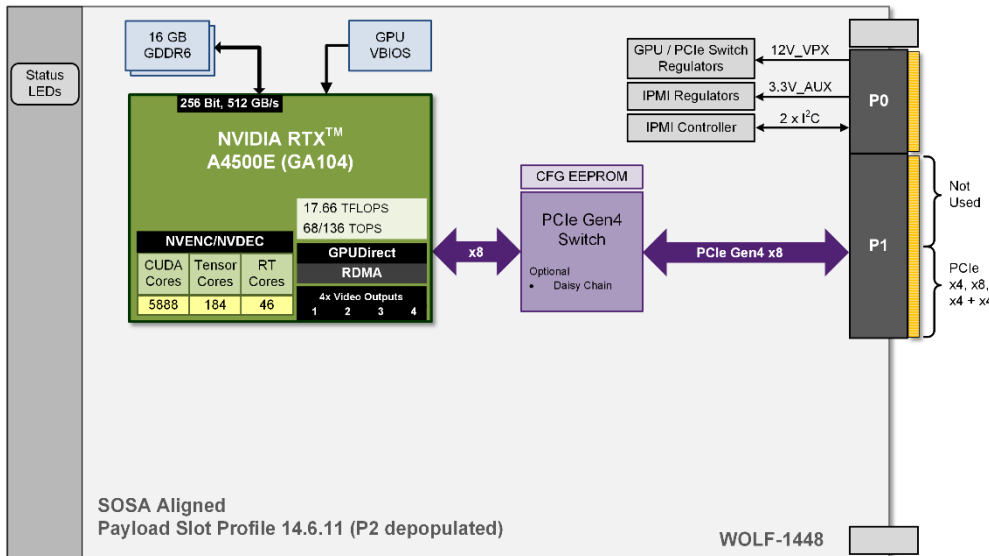
This information is subject to change

IN DEVELOPMENT



Standard model with video output over P2. A configurable PCIe Gen4 switch provides support for OpenVPX slot profiles.

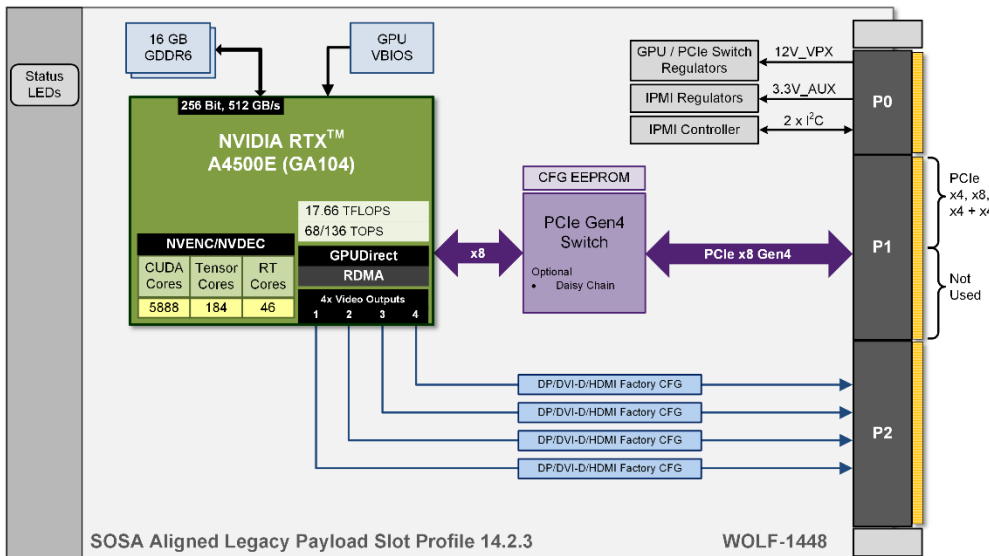
PCIe x4, x8, x16
x4 + x4,
x8 + x4 + x4,
x8 + x8



SOSA aligned model with support for Payload Slot Profile 14.6.11-0. The P2 connector is depopulated for this model.

A configurable PCIe Gen4 switch provides support for up to PCIe x8 over P1B.

Not Used
PCIe x4, x8, x4 + x4



SOSA aligned model with support for Legacy Payload Slot Profile 14.2.3. The P2 connector is populated and can be used for video output.

A configurable PCIe Gen4 switch provides support for up to PCIe x8 over P1A.

PCIe x4, x8, x4 + x4
Not Used

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.

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 codecs, including 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.

CONFIGURABLE PCIe GEN4 INTERFACE

Ampere is the first NVIDIA GPU generation to include support for PCIe Gen4, providing double the throughput of the previous generation. The WOLF-1448 module also includes a configurable PCIe Gen4 switch with Non-Transparent Bridging (NTB) and daisy chain options.

WOLF ADVANCED METAL

WOLF has invested significant time and resources into developing a superior cooling material for VPX module heat plates. This material provides lower thermal resistance between the GPU and the chassis cold wall, allowing the GPU to run at higher power levels without risking thermal throttling. This material also reduces the overall weight of the unit, allowing system integrators to meet their SWaP requirements.

OPENVPX AND SOSA SLOT PROFILE SUPPORT

This module's configurable PCIe interface provides support for several OpenVPX slot profiles. The module ICD is pin compatible with older generation WOLF-1348 and 1116 modules, enabling a simple plug-in upgrade.

The module can also be configured to be SOSA aligned with the following SOSA profiles supported:

- SLT3-PAY-1F1U1S1S1U1U2F1H-14.6.11-0 Payload Slot Profile, P2 depopulated
- SLT3-PAY-1F1U1S1S1U1U4F1J-14.6.13-0 Payload Slot Profile, P2 depopulated
- SLT3-PAY-2F2U-14.2.3 Legacy Payload Slot Profile

This information is subject to change

ORDERING CODES

The following table defines series of common order codes for the VPX3U-A4500E-VO module. The asterisks denote characters of the part number that are defined based on common configuration options. Some configuration options for this module include:

- Display Interfaces
- Variant Locked
- P2 depopulated (SOSA Aligned)
- Conformal Coatings
- Default Power Threshold
- PCIe Configuration Options

Ordering Number	Description
3U VPX Ampere A4500 Single Slot Configurations	
144833-F001-***VPX3vA0	3U VPX, Conduction Cooled, 1", NVIDIA Ampere A4500, Rear IO: 4x DP, OpenVPX
144833-F005-***VPX3vA0	3U VPX, Conduction Cooled, 1", NVIDIA Ampere A4500, P2 depopulated, SOSA Payload (no video IO)
144833-F005-***VPX3vA0	3U VPX, Conduction Cooled, 1", NVIDIA Ampere A4500, Rear IO: 4x DP, SOSA Legacy

* 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)



This information is subject to change

Datasheet Rev.11