

VPX3-4924

3U VPX 6.2 TFLOPs GPGPU Processor Card with
NVIDIA Tesla Pascal P6 HPC 16 nm GPU

**CURTISS-
WRIGHT**

CURTISSWRIGHTDS.COM



Key Features

- NVIDIA Tesla P6 16 nm GPU
- 2048 NVIDIA CUDA® cores
- 6168 GFLOPS single precision (FP32)
- 67 GFLOPS/watt
- 16 GB GDDR5 with NVIDIA GPUDirect™ DMA technology
- Memory Width: 256 bit
- Max memory bandwidth: 192 GB/s
- NVIDIA GRID™ vGPU virtualization
- PCIe x16 Gen 3

Applications

- ISR and EW applications where TFLOPs of accelerated processing is required
- Massive data ingest of modern Radar, SIGINT, EO/IR sensors
- Unparalleled HPEC performance in cross-cueing applications

Overview

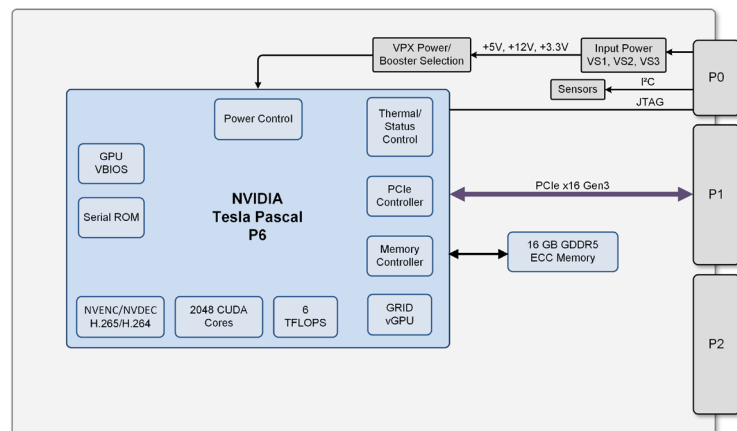
The rugged [VPX3-4924](#) NVIDIA® Tesla® Pascal GPU processor board is designed and manufactured by Wolf Advanced Technology. This board is part of a family of GPGPU modules available from Curtiss-Wright Defense Solutions to enable development of High Performance Embedded Computing (HPEC) systems.

Leveraging the NVIDIA® Tesla® Pascal 16 nm GPU technology to bring extreme high performance, the VPX3-4924 brings GPUDirect™ DMA and GRID™ virtualization technology to ruggedized embedded platforms. NVIDIA GRID™ is the industry's most advanced technology for sharing true virtual GPU (GRID 2.0 vGPU™) hardware acceleration between multiple users; max of 16 instances. This technology ensures complete application compatibility, which means features and experiences are the same as they would be on a physical device.

The NVIDIA High Performance Compute (HPC) mode makes the VPX3-4924 especially suitable for use as an advanced GPGPU compute engine. Starting with Pascal generation boards, switching between graphics and compute mode is no longer required. The NVIDIA GRID automatically handles setting the adjustments. The HPC mode exposes a large 16 GB Base Address Register (BAR) for direct access to the frame buffer from the CPU and other PCI Express® (PCIe) devices.

Tesla Pascal GPU architecture also provides a more powerful Unified Memory feature. Pascal's larger virtual memory address space enables GPUs to access the entire system memory plus the memory of all GPUs in the system. The on-demand page migration engine allows the system to migrate pages from anywhere in the system to the GPU's memory for efficient processing. This improved memory handling results in significantly improved algorithm efficiency.

Due to the critical importance of size, weight, and power (SWaP) in aerospace and defense applications, the GPU on the ruggedized VPX3-4924 can be tuned to maximize GPGPU capability while minimizing power usage.



VPX3-4924 block diagram

WOLF
ADVANCED TECHNOLOGY

TRUSTED
PROVEN
LEADER

INFO: CURTISSWRIGHTDS.COM
EMAIL: DS@CURTISSWRIGHT.COM

Specifications and Features

Processor

- NVIDIA Tesla Pascal P6
 - + 2048 CUDA cores, 6.2 TFLOPS
 - + 16 GB GDDR5 ECC memory
 - + Memory width: 256-bit
 - + Maximum memory bandwidth: 192 GB/s
- Host interface: x16 PCIe
- NVENC/NVDEC accelerator for HVEC (H.265) and AVC (H.264) hardware encode/decode

Special Features

- High Performance Compute (HPC) mode
- NVIDIA GRID™ vGPU virtualization with up to 16 vGPU instances
- Full 16 GB DMA BAR access

Power

- +5 (VS3) or +12V (VS1) centric, +3.3V_AUX also required
- Configurable GPU power cap: 40-90W
- Maximum operating power: 90W

Environmental

- Operating temperature (air-cooled): -20 to 65°C
- Operating temperature (conduction-cooled): -40 to 65°C
- Humiseal 1B73 conformal coating
- Other environmental specifications are per [Wolf Advanced Technology](#)

Software Support

- NVIDIA drivers supporting Linux® (minimum requirements)
 - + CUDA Toolkit 9, CUDA Compute version 6.1
 - + OpenCL™, DirectX® 12, OpenGL 4.5
 - + Vulkan 1.0

Ordering Information

TABLE 1		VPX3-4924 ordering information
PART NUMBER	VARIANTS	
VPX3-4924-A141-000	3U OpenVPX module with NVIDIA Tesla P6 <ul style="list-style-type: none"> > 2048 CUDA cores, 16 GB GDDR5, 6.2 TFLOPS > Air-cooled, “1.0” pitch, temperature range (-20 to 65°C) > 16 lane PCIe Gen3 > Configurable power 40-90W > 5V centric, no display outputs 	
VPX3-4924-A141-001	3U OpenVPX module with NVIDIA Tesla P6 <ul style="list-style-type: none"> > 2048 CUDA cores, 16 GB GDDR5, 6.2 TFLOPS > Air-cooled, “1.0” pitch, temperature range (-20 to 65°C) > 16 lane PCIe Gen3 > Configurable power 40-90W > 12V centric, no display outputs 	
VPX3-4924-C141-000	3U OpenVPX module with NVIDIA Tesla P6 <ul style="list-style-type: none"> > 2048 CUDA cores, 16 GB GDDR5, 6.2 TFLOPS > Conduction-cooled, “1.0” pitch, temperature range (-40 to 65°C) > 16 lane PCIe Gen3 > Configurable power 40-90W > 5V centric, no display outputs 	
VPX3-4924-C141-001	3U OpenVPX module with NVIDIA Tesla P6 <ul style="list-style-type: none"> > 2048 CUDA cores, 16 GB GDDR5, 6.2 TFLOPS > Conduction-cooled, “1.0” pitch, temperature range (-40 to 65°C) > 16 lane PCIe Gen3 > Configurable power 40-90W > 12V centric, no display outputs 	