



NEWS RELEASE

FOR IMMEDIATE RELEASE

Contact: John Wranovics
M: 925.640.6402
jwranovics@curtisswright.com

Curtiss-Wright Adds New NVIDIA® Tesla® Pascal™ 16 nm GPGPU Processor Modules to its Family of HPEC Processors for Demanding ISR/EW Applications

Tesla P6-based VPX3-4924 3U OpenVPX™ (single GPU) and VPX6-4944 6U OpenVPX™ (dual GPU) Modules enhance virtualization and data transfers in HPEC systems

ASHBURN, Va. – October 1, 2018 – Curtiss-Wright's Defense Solutions division today announced that it has expanded its family of [open architecture high performance embedded computing \(HPEC\) processors](#) designed for demanding ISR applications with the addition of two new NVIDIA® Tesla® Pascal™ (P6) 16 nm GPGPU-based OpenVPX modules. Designed for use in compute intensive ISR and EW systems, the fully rugged VPX3-4924 and VPX6-4944 modules are ideal for use in applications, such as sensor fusion or large swath video mapping, that require TFLOPS of accelerated processing. With their on-device support for NVIDIA High Performance Compute (HPC) mode, these HPEC engines can ingest the massive amounts of data generated by modern Radar, SIGINT, and EO/IR sensors. The [VPX3-4924 module, a 3U OpenVPX GPGPU Processor](#), features a single Tesla P6 GPU and delivers 6.2 TFLOPS performance. For more demanding applications, the [6U form factor VPX6-4944](#) deploys dual Tesla P6s, doubling available compute power to 12.4 TFLOPS. These [state-of-the-technology GPGPU modules](#) extend Curtiss-Wright's proven leadership as a supplier of the most advanced compute solutions for embedded ISR applications.

These GPGPU modules enable system designers to fully leverage the benefits of virtualization via NVIDIA's Virtual GPU (vGPU™) graphics virtualization platform, and provide virtual machines (VM) access to NVIDIA GPU technology. vGPU is supported by NVIDIA's GRID™ 2.0 API, the industry's most advanced technology for sharing virtual GPU hardware acceleration between multiple users, with support for up to 16 instances. Using the GRID API, virtual and physical devices present the system designer with identical features and an identical user experience, making the actual and the virtual GPU instance indistinguishable.

NVIDIA High Performance Compute (HPC) Mode

Available only on the Tesla line of NVIDIA GPUs, HPC mode is especially well suited for use as an advanced GPGPU compute engine. Starting with Pascal-generation boards, switching between graphics and compute mode is no longer required. NVIDIA GRID software automatically handles

most setting adjustments previously requiring the switch to graphics mode. In HPC mode, the Tesla Pascal includes ECC (error-correcting code) memory support.

“With these latest additions to our family of NVIDIA Pascal-based modules, Curtiss-Wright continues our commitment to being the proven and trusted HPEC solution provider, enabling our ISR/EW customers to achieve their parallelization objectives and maximize their performance per watt”, said Lynn Bamford, Senior Vice President and General Manager, Defense Solution division. “The Tesla P6-based VPX3-4924 and VPX6-4944 GPGPU modules, with their large 16 GB DMA BAR memory per GPU, deliver unmatched communications in support of I/O intensive HPEC applications such as sensor fusion or large swath video mapping. Even better, our OpenHPEC Accelerator Suite of software tools reduces program risk for developing complex GPGPU HPEC systems, greatly expediting the time between development and deployment.”

Key Performance Features per GPU:

- 2048 NVIDIA CUDA® cores
- 6.2 TFLOPS single precision (FP32)
- 67 GFLOPS/Watt
- 16 GB GDDR5 Memory
- NVIDIA GPUDirect™ DMA technology
- Memory Width: 256 bit
- Max memory bandwidth: 192 GB/s
- ECC memory support
- NVIDIA GRID vGPU virtualization
- HPC Mode: 16 DMA BAR
- PCIe x16 Gen 3

Complete System-Level HPEC Solutions

The VPX3-4924 and VPX6-4944 modules are fully interoperable with Curtiss-Wright's [broad family of system-level C4ISR/EW OpenVPX solutions](#), such as the [Xilinx® FPGA-based VPX3-534/535 3U modules](#), and the [Intel® Xeon®-D processor-based CHAMP-XD1 single processor 3U](#) and [CHAMP-XD2 dual processor 6U DSP modules](#). Numerous system configurations can be formed using combinations of these boards. For example, a single CHAMP-XD2 node can be used to control both Tesla P6 GPUs on a VPX6-4944, or each node on a CHAMP-XD2 can be mapped to each individual GPU. Taking this design one step further, a single Xeon-D on the CHAMP-XD2 can control one VPX6-4944 board upstream, while the other Xeon-D maps to a second VPX6-4944 module downstream. This processing slice, consisting of a CHAMP-XD2 and two VPX6-4944 modules, delivers an impressive 25 TFLOPS of performance that can be connected to the rest of the system via 40 GB Ethernet/InfiniBand ports on the OpenVPX data plane.

OpenHPEC™ Accelerator Suite™ Support

Curtiss-Wright HPEC modules and systems are supported by the [OpenHPEC Accelerator Suite](#) of best-in-class software development tools. These powerful tools enable system developers to achieve their software implementation faster. For example, the suite includes the powerful Bright Cluster Manager from [Bright Computing](#), an NVIDIA partner. Bright Cluster Manager provisions and

monitors both the CPU and GPU boards, and includes a fully configurable module environment. The OpenHPEC tool suite includes Bright's deep learning libraries and tools from both Intel and NVIDIA, including Caffe and TensorFlow. It also provides the Arm Forge which enables true system level debugging and profiling for both CPU and GPUs, and supports MPI, OpenMP and ACC.

For high-speed, low-latency, peer-to-peer communications, the OpenHPEC Accelerator Suite also includes Dolphin's PCIe communication library, which hides the complexities of directly programming the system's PCIe devices. In addition to supporting GPU sharing between the CPUs, the Dolphin library also supports both CPU direct and remote direct memory access (RDMA).

Use of the OpenHPEC Accelerator Suite simplifies, speeds, and lowers the cost of ISR application development. These tools deliver the benefits of open standard High Performance Computing (HPC) software to the COTS market to effectively remove risk when developing large scale embedded computer clusters.

Resulting from its Reseller Agreement with [WOLF Advanced Technology](#), the high-performance VPX3-4924 and VPX6-4944 HPEC modules have been pre-validated. They complement Curtiss-Wright's previously announced family of NVIDIA Pascal GPGPU modules by speeding and easing the integration of HPEC solutions into deployed systems.

For more information about the Curtiss-Wrights Defense Solutions division, please visit www.curtisswrightds.com.

About Curtiss-Wright Corporation

Curtiss-Wright Corporation is a global innovative company that delivers highly engineered, critical function products and services to the commercial, industrial, defense and energy markets. Building on the heritage of Glenn Curtiss and the Wright brothers, Curtiss-Wright has a long tradition of providing reliable solutions through trusted customer relationships. The company employs approximately 8,600 people worldwide. For more information, visit www.curtisswright.com.

###

Note: All trademarks are property of their respective owners.