NEWS RELEASE

Curtiss-Wright Announces New NVIDIA® Quadro® Pascal™ 5200 GPGPU Processor Modules for ISR/EW and AI Applications

VPX3-4933 3U OPENVPX™ (SINGLE GPU) AND VPX6-4953 6U OPENVPX™ (DUAL GPU) MODULES ENHANCE VIRTUALIZATION AND DATA TRANSFERS IN HPEC SYSTEMS

ASHBURN, Va. – June 27, 2019 – Curtiss-Wright's Defense Solutions division, a trusted leading supplier of rugged ISR and EW processing modules and systems, today announced that, through its Reseller Agreement with WOLF Advanced Technology, 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 Quadro Pascal (P5200) GPGPU/Inference Engine-based OpenVPX modules. Designed to support compute-intensive ISR and EW systems, and artificial intelligence (AI) applications, the fully rugged VPX3-4933 and VPX6-4953 modules are ideal for use in deployed ISR and EW applications (including deep learning) where TFLOPs of accelerated processing are required. These applications include high-performance radar, SIGINT, EO/IR, data fusion ingest, processing, and display.

The VPX3-4933 module, a 3U OpenVPX GPGPU processor, features a single NVIDIA Quadro P5200 GPU that delivers 8.7 TFLOPS performance. For more demanding applications, the 6U form factor VPX6-4953 hosts dual P5200s, doubling available compute power to 17.4 TFLOPS. These state-of-the-art GPGPU modules further extend Curtiss-Wright's proven leadership as a supplier of the most advanced computing solutions for embedded ISR applications.

To meet demanding rugged military and aerospace specifications, both GPGPU boards feature a chip-down design. The P5200 GPU, in addition to providing top-of-the-line processing power

from its 2560 CUDA® cores, also provides a maximum memory bandwidth of 243 GB/s, the largest available in the embedded GPU market. Due to the critical importance of size, weight, and power in aerospace and defense applications, the GPU(s) on these ruggedized boards is tune-able to maximize GPGPU capability while minimizing power usage.

"With the introduction of these new GPGPU modules, Curtiss-Wright enhances its ability to support customers designing deployed artificial intelligence applications," said Lynn Bamford, Senior Vice President and General Manager, Defense and Power. "These latest additions to our family of NVIDIA Pascal-based modules also further strengthen our commitment to being the embedded industry's proven and trusted HPEC solution provider."

VPX3-4933 (3U/Single GPU) Performance Features:

- 3U VPX GPGPU Processor Card with Chip-Down NVIDIA Quadro Pascal P5200
- 8.7 TFLOPS/2560 CUDA cores
- 4 simultaneous video outputs: 3 DisplayPort++ 1.4 & 1 DVI
- 16GB GDDR5 256-bit Memory: 243 GB/s max. bandwidth to each GPU
- PCIe Gen 3 x8 or x16
- Operating power configurable hard cap: 40-150W

VPX6-4953 (6U/Dual GPU) Performance Features:

- 6U VPX GPGPU Processor Card with Chip-Down Dual NVIDIA Quadro Pascal P5200s
- 17.4 TFLOPS/5120 CUDA cores
- 8 simultaneous video outputs: 3 DisplayPort++ 1.4 & 1 DVI
- 32GB GDDR5 256-bit Memory: 243 GB/s max. bandwidth to each GPU
- PCIe Gen 3 x16
- Operating power configurable hard cap: 100-260W

Key Performance Features per NVIDIA P5200 GPU:

- 2560 NVIDIA CUDA cores
- 8.7 TFLOPS single precision (FP32)
- 67 GFLOPS/Watt
- 16 GB GDDR5 Memory
- NVIDIA GPUDirect™ DMA technology

- Memory Width: 256-bit
- Max memory bandwidth: 243 GB/s
- PCIe x16 Gen 3

Complete System-Level HPEC Solutions

The VPX3-4933 and VPX6-4953 modules are fully interoperable with Curtiss-Wright's broad family of 3U and 6U system-level C4ISR/EW OpenVPX solutions. SWAP-constrained systems can pair the single GPGPU VPX3-4933 with our Intel® Xeon®-D processor-based CHAMP-XD1 DSP engine, which together can augment the powerful sensor processing capabilities of our Xilinx® FPGA-based VPX3-534/535 transceiver modules. For higher performance 6U systems, the dual GPGPU VPX6-4953 can pair with our dual Intel Intel® Xeon®-D processor based CHAMP-XD2 DSP module and powerful CHAMP-FX4 triple-FPGA processer module workhorse, along with many other modules in our Fabric40™ family of 40Gbps products. Numerous system configurations can be formed using combinations of these boards. For example, a single CHAMP-XD2 node can be used to control both Pascal P5200 GPUs on a VPX6-4953, 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 VPX6-4953 board upstream, while the other Xeon-D maps to a second VPX6-4953 module downstream. This processing slice, consisting of a CHAMP-XD2 and two VPX6-4953 modules, delivers an impressive 36 TFLOPS of performance that can be connected to the rest of the system via multiple 40 Gbps Ethernet/InfiniBand™ ports available on the OpenVPX data plane, and through a Curtiss-Wright VPX6-6802 central fabric switch.

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 develop their software 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 suite, 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 VPX3-4933 and VPX6-4953 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 on Curtiss-Wright Defense Solutions products, 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 9,000 people worldwide. For more information, visit www.curtisswright.com.

###

NOTE: All trademarks are property of their respective owners.