


## EnsembleSeries™ GSC6202

6U OpenVPX GPGPU accelerator powered by dual NVIDIA Pascal or Maxwell processors

- NVIDIA Pascal or Maxwell GPGPU processing
- Industry-standard GPU MXM cards for upgradability
- 5.8-12.8 TFLOPS, 16-32GB of GDDR5 and PCIe Gen 3 interconnects
- Application development with OpenCL, NVIDIA CUDA
- Advanced mechanical packaging enables successful rugged deployment of MXMs



The EnsembleSeries™ GSC6202 accelerator is a 6U OpenVPX carrier module that integrates two high-performance NVIDIA GPGPUs for applications that can benefit from massively parallel processing on streams of high-bandwidth data delivering some of the highest GFLOP processing performance as well as the highest GFLOP/Watt performance efficiency in the industry.

### Meeting Today's Processing Demands

Commercial and defense applications such as radar, electro-optical/infrared (EO/IR), artificial intelligence (AI), electronic warfare (EW) and sensor fusion applications generate large amounts of raw sensor or network data that need to be processed in real-time to extract actionable intelligence. Each new generation of sensor arrays ushers in higher resolutions and frame rates. By offloading compute-intensive operations to GPGPUs such as fast Fourier transforms (FFTs), matrix multiplication, constant false alarm rate (CFAR), QR decomposition (QRD), synthetic aperture radar (SAR), video codecs (H.264, JPEG2000), pattern recognition or deep packet inspection, system architects can engineer solutions that can meet today's processing demands — with room to scale for higher

performance requirements in the future while preserving significant IP investment. The Pascal architecture solution also provides accelerated neural network and artificial intelligence performance for deep learning and cognitive and adaptive algorithms.

The EnsembleSeries GSC6202 accelerators are Mercury's 4th generation 6U dual GPGPU module, backed with a history of field-proven deployments in a number of defense programs since 2008. Updates to the EnsembleSeries GSC6202 include increased memory to 8GB-16GB per GPU (16-32GB per module), higher I/O bandwidth via PCIe Gen3 and support for NVIDIA Pascal P5000 or Maxwell M6 MXM GPUs.

### Pascal P5000 or Maxwell M6 GPGPU

The EnsembleSeries GSC6202 accelerators are powered by NVIDIA's next generation professional graphics cards, the Pascal P5000 or Maxwell M6 GPU, to the embedded processing environment. The new devices take advantage of NVIDIA's advancing MXM architectures and their built-in graphics processor to provide increased performance and memory bandwidth. The M6 features 1536 shading units (cores) at up to 930 MHz and can eclipse 2.9 TFLOPS of single precision floating point. 16GB of GDDR5 using a 256-bit memory interface provides up to 160.4

*Mercury Systems is the better alternative for affordable, secure processing subsystems designed and made in the USA. These capabilities make us the first commercially based defense electronics company built to meet rapidly evolving next-generation defense challenges.*



ACQUIRE



DIGITIZE



PROCESS



STORAGE



EXPLOIT



DISSEMINATE

GB/s of memory bandwidth per M6 GPU. The P5000 MXM delivers 2048 shading units, 6.4 TFLOPS (single precision), and 16 GB GDDR5 per GPU (32 GB per GSC6202) for 192 GB/s of memory bandwidth. In both cases, the PCIe 3.0 x16 interface provided by the MXM 3.1 card provides high throughput for streaming data both to and from the GPU memory.

## EnsembleSeries GSC6202 Advancements

To keep up with the rapid evolution of GPGPU technology, the EnsembleSeries GSC6202 accelerators are carrier cards built using two industry-standard mobile PCIe modules (MXM) from NVIDIA. Because of the modular nature of the MXM, the EnsembleSeries GSC6202 carrier card only needs to be engineered once while maintaining the flexibility of updating GPUs. This approach saves significant engineering development time, allowing programs to deploy with the latest, highest performing embedded GPGPU technology as soon as they are available from NVIDIA. The option to configure the GSC6202 with either M6 or P5000 MXMs demonstrates the value of this approach.

The EnsembleSeries GSC6202 improves upon the previous generation modules by doubling the interconnect bandwidth via a PCIe Gen 3 switch. By adding the required throughput and MXM 3.1 support, this module can support the leading edge NVIDIA GPGPUs described above.

## Open Software Environment

Mercury leverages over 35 years of multicomputer software expertise across its many platforms, including the latest multicore processors found in GPGPUs. This strategy is fully applied to the EnsembleSeries GSC6202 accelerators. Because the processor, memory and surrounding technologies are leveraged across product lines, software developed on the GSC6202 can interface seamlessly with other Mercury products.

EnsembleSeries GSC6202 accelerators interface with Intel modules running Red Hat® Linux®. Several software development environments are available for GSC6202 accelerators:

- NVIDIA CUDA: A parallel computing architecture that is accessible to software developers through industry standard programming languages.
- OpenCL: An open-source standard for cross-platform and parallel programming.

## System Management

EnsembleSeries GSC6202 accelerators implements the advanced system management functionality architected in the VITA 46.11 specification to enable remote monitoring, event management, and hardware revision and health status. Using the standard I2C bus and intelligent platform management controller (IPMC) protocol, the on-board system management block implements the IPMC.

This allows the EnsembleSeries GSC6202 accelerators to:

- Read sensor values
- Read and write sensor thresholds, allowing an application to react to thermal, voltage or current variations that exceed those thresholds
- Reset the entire module
- Power up/down the entire module
- Retrieve module field replaceable unit (FRU) information
- Be managed remotely by a chassis management controller at the system level, such as implemented on EnsembleSeries 6U OpenVPX switches

## VPX-REDI

The VPX (VITA 46) standard defines 6U boards with a modern, high-performance connector capable of supporting today's high-speed fabric interfaces. VPX is most attractive when paired with the ruggedized enhanced design Implementation standard – REDI (VITA 48). EnsembleSeries GSC6202 accelerators are implemented as 6U VPX-REDI conduction-cooled (VITA 48.2), Air Flow-By™ (VITA 48.7), or Liquid Flow-Through (VITA 48.4) with an air-cooled variant available in the same VPX form-factor for less rugged environments.

Targeted primarily for harsh-environment embedded applications, VPX-REDI offers extended mechanical configurations supporting higher functional density, such as two-level maintenance (2LM). 2LM allows maintenance personnel to replace a failed module and restore the system to an operational state in a limited time period, minimizing potential damage to the module.

## Technical Specifications

### GPGPU

Two NVIDIA GPGPU MXMs designed specifically for embedded GPGPU applications

Option 1: Dual NVIDIA P5000 Pascal architecture MXM

- 4096 total processing cores (2048 per MXM)

- 12.8 peak theoretical single-precision TFLOPS (6.4 peak theoretical TFLOPS per MXM)

- x32 total PCIe 3.0 lanes (x16 PCIe 3.0 per MXM)

- 32 GB total GDDR5 Memory (16 GB per MXM)

- 256-bit memory interface

- 384 GB/s memory bandwidth (192 GB/s per MXM)

Option 2: Dual NVIDIA M6 Maxwell architecture MXM

- 3072 total processing cores (1536 cores per MXM)

- 5.8 peak theoretical single-precision TFLOPS (2.9 peak theoretical TFLOPS per MXM)

- x32 total PCIe 3.0 lanes (x16 PCIe 3.0 per MXM)

- 16 GB total GDDR5 memory (8 GB per MXM)

- 256-bit memory interface

- 320 GB/s memory bandwidth (160 GB/s per MXM)

4 Display Port display outputs (2 outputs per MXM)

- To front-panel (air-cooled only) and OpenVPX backplane

2 analog VGA display outputs (one per MXM)

- To front-panel (air-cooled only) and OpenVPX backplane

### 64-Lane Configurable PCIe Switch

Configurable switch allows for multiple system-level configurations in terms of non-transparent bridging and enumeration

x16 PCIe 3.0 connections to each MXM site (32 lanes total)

x32 PCIe 3.0 total connections to backplane

- x16 PCIe 3.0 OpenVPX P2 expansion plane

- x16 PCIe 3.0 OpenVPX P5 expansion plane

### IPMI (System Management)

On-board IPMI controller

Voltage and temperature monitor

Geographical address monitor

Power/reset control

On-board CPLD, FRU EEPROM interfaces

### OpenVPX Multi-Plane Architecture

System management via IPMB-A and IPMB-B link on P0 management plane

Dual full x16 or dual x8 PCIe on P2 and P5 expansion plane

4 DisplayPort display outputs on P6 mezzanine I/O plane

2 analog VGA outputs on P3 mezzanine I/O plane

### Mechanical

6U OpenVPX (air-cooled, Air Flow-By, conduction-cooled, or Liquid Flow Through)

1.0" pitch, single-slot

OpenVPX and VPX-REDI


### Power Consumption

Typically 140W per accelerator

### Environmentals

For information on cooling and rugged packaging options, see our [MOTS+ extreme environmental protections](#)

## GSC6202 Functional Block Diagram



The Mercury Systems logo is a registered trademark of Mercury Systems, Inc. Other marks used herein may be trademarks or registered trademarks of their respective holders. Mercury products identified in this document conform with the specifications and standards described herein. Conformance to any such standards is based solely on Mercury's internal processes and methods. The information contained in this document is subject to change at any time without notice.

Copyright © 2019 Mercury Systems, Inc.

3283.05E-0124-ds-GSC6202



INNOVATION THAT MATTERS®

### CORPORATE HEADQUARTERS

50 Minuteman Road • Andover, MA 01810 USA  
(978) 967-1401 • (866) 627-6951 • Fax (978) 256-3599

### EUROPE MERCURY SYSTEMS, LTD

Unit 1 - Easter Park, Benyon Road, Silchester, Reading  
RG7 2PQ United Kingdom  
+ 44 0 1189 702050 • Fax + 44 0 1189 702321

