



GENERAL MICRO SYSTEMS, INC.
TRUSTED AND DEPLOYED SINCE 1979



LIGHTNING X422

Deep Learning, Artificial Intelligence, Rugged, Conduction-Cooled, Fully Sealed GPGPU Co-Processor

Dual
220 TFLOPS
GPGPU Processors

32 Lanes, 8 Gbps
3.0 PCIe
Dual Bus Expansion

Flow-Through
FlexIO™
Daisy Chain Chassis

Operating Temperature Range
-20° to 55°C
Available in Rugged Level 1-2

SYSTEM HIGHLIGHTS

- PCIe-over-cable local bus extension for inter-/intra chassis expansion for co-processors
- Supports up to two full size commercially-available General-Purpose Computing on Graphics Processing Units (GPGPUs)
- Accommodates specialized FPGA, ASIC, co-processor cards
- Conduction-cooled rugged enclosure provides 600 W of power
- FlexIO™ backplane provides PCIe Gen 3 at 8 Gbps
- Independent GPGPU outputs
- GPGPU modules support a total of 32 PCIe Gen 3 lanes, 16 lanes of input and 16 lanes of output
- X422 system daisy chains to multiple GPGPU co-processor units
- Software support for NVIDIA™ CUDA proprietary framework and OpenCL open framework

SYSTEM OVERVIEW

The X422 is a GPGPU Co-Processor System that supports GMS' companion S422 SW Server. Up to two full size GPGPU cards can be installed, either as two independent processing units or they can be intelligently linked together to form one virtual processing unit. In addition to the GPGPU capabilities, it offers bus expansion capabilities to virtually any full size, full height PCIe 3.0 card, but is optimized to harness multiple GPGPU systems. The X422 affords the use of commercially available, full size, GPGPU products in a completely sealed and protected, fully ruggedized enclosure that utilizes GMS' patented RuggedCool™ technology.

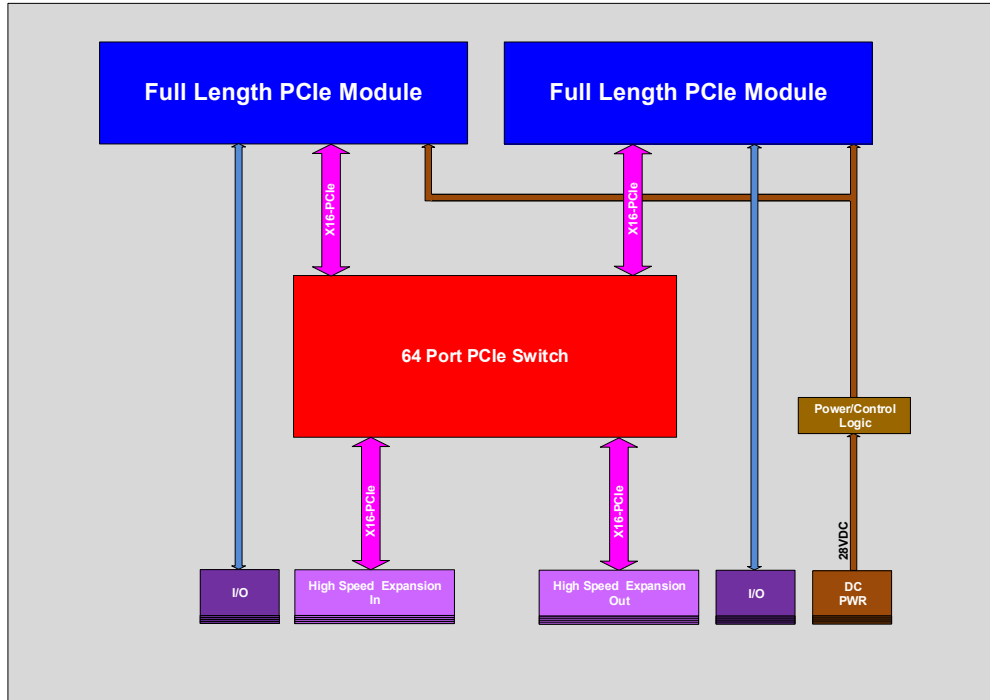
Owing to the underlying GPGPU architecture with multiple cores and threads, the X422 can accelerate large data processing tasks such as image recognition, digital signal processing, data mining, block chain computation, artificial intelligent, machine vision, image processing, vector processing, and other computing-intensive tasks. PCIe 3.0 bus expansion allows upward-scaling by cascading (daisy chaining) additional GPGPUs as necessary.

The X422 utilizes GMS's FlexIO™ flow-through architecture which is based upon wide PCIe 3.0 lanes (x16) operating at 8 Gbps. When connected via x16 PCIe 3.0 to the companion S422 dual Xeon® conduction-cooled rugged server, the GPGPUs (or any PCIe cards) appear to be within the same "bus" of the main server. This closely-coupled architecture allows for rapid data passing, RDMA and "atomic" operations, or provides a fully autonomous GPGPU co-processor where only data is shared asymmetrically from the main CPUs.

The onboard intelligent PCIe 3.0 switch provides packet processing and local "routing", allowing the user to customize a homo- or heterogeneous architecture between the local GPGPU resources. For example, the X422 can be partitioned for GPGPU A to operate separately and independently from GPGPU B—or they can pass data between each other via the dual x16 PCIe 3.0 fabric. Separate, GPGPU-specific, and customizable I/O "pipes" are available to the rear panel for each PCIe slot. This allows sensor-to-GPGPU processing, down- or upstream processing, or an additional data path that bypasses the companion S422 dual-Xeon server.



BLOCK DIAGRAM



AVAILABLE NVIDIA® GPGPUs

The X422 supports the following NVIDIA® GPGPU cards.

(Model / Architecture)	NVIDIA CUDA® Cores	Single Precision GFLOPS	TDP (W)
Jetson TX2 / Pascal	256	1500	8
GTX 1060 / Pascal	1280	3594	80
GTX 1080 / Pascal	2560	7967	150
P4 PCIe – Tesla / Pascal	2560	4800	67
P100 PCIe – Tesla / Pascal	3584	10000	250
TITAN Xp / Pascal	3840	10790	250
V100 PCIe – Tesla / Volta	5120	14026	250
TITAN V - Volta	5120	12288	250
TITAN V – Volta [Tensor]	5120	110592	250

ENVIRONMENTAL SPECS

- SWaP-E:
- Size: 12.45" x 11.6" x 2.54" (including fins)
- Weight:
- Power: 300 W (one card) to 550 W (two cards)
- MIL-STD: MIL-STD-810G, MIL-STD-1275D, MIL-S-901D, MIL-STD-461F, DO-160D, IP64
- Temperature: Operates up to extended temp -20°C to +55°C
- Ruggedness: Available in ruggedization levels R1-R2



APPLICATIONS

(The X422 is a high performance GPGPU co-processor and bus expansion system designed for use with the S422 SW Server.

The X422 is an ideal forwardly-deployed, vehicle-mounted, high performance computing system. Applications include computing clusters and parallel computing, digital signal processing, digital image processing, video processing, neural networks, data mining, cryptography, and intrusion detection.

The X422 system is fully compliant to MIL-STD-810G, MIL-STD-1275D, MIL-S-901D, DO-160D, MIL-STD-461E and has ingress protection up to IP64.

I/O AND EXPANSION OPTIONS

- PCIe-over-cable bus extension for co-processor cards
- Up to two full size GPGPUs (dual x16 PCIe 3.0 slots)
- 32 lanes PCIe Gen 3, dual x16
- Two independent card-specific fabric or I/O ports (one per PCI slot). Example: Titan V provides 2x DisplayPort and 1x HDMI (A), and 3x DisplayPort and 1x HDMI (B)

RUGGEDIZATION LEVELS

	TEMP	SHOCK	VIBRATION	MAX IP LEVEL
RUGGED 1	0° - 55°C	20G	.0001 g ² /Hz	54
RUGGED 2	-20° - 55°C	20G	.0008 g ² /Hz	64

* Vibration frequency for systems tested between 5Hz – 2000Hz