



NVIDIA® GPGPU Accelerator Pack

Accelerated Integration and Deployment


Designed to simplify at-scale NVIDIA GPUs integration and deployment (up to 255 GPGPUs per platform), AMI's GPGPU Accelerator Pack delivers a seamless management and monitoring experience through stable, industry-standard BMC features for modern-day server platforms designed to accelerate datacenter and cloud workloads.

Save Time, Remain Stable, Keep Updating

AMI's GPGPU Accelerator Pack is an advanced technology add-on package designed to extend the functionality of the BMC to the GPGPU Accelerator board. Available on AMI's MegaRAC® platform management codebase, customers can experience a comprehensive out-of-band (OOB) manageability solution validated by industry-specification compliance tools, AMI internal tests, Arm® (SBMR), and NVIDIA (NVVS) validation suites. Customers are also able to contact AMI's global support teams for expert assistance, issue resolution, and customization requests.

Key Features and Benefits

- ✓ Out-of-the-box, Enterprise-ready solution
- ✓ Stable, hardened, validated code-base
- ✓ Complete Inventory of subsystem components
- ✓ Advanced add-on technologies (TPs/EPs)
- ✓ Integration with AMI UEFI & HRoT products
- ✓ Easily Scalable across GPU clusters
- ✓ Multi-generation NVIDIA GPU support
- ✓ NVIDIA Validation Suite (NVVS) compliant
- ✓ Extensive Power Management
- ✓ Fast, configurable FW Updates & Resiliency
- ✓ Secure transactions and communication within the GPU complex and external components



AI-Generated Image



Comprehensive Accelerator Management

AMI's industry-leading GPGPU Accelerator Pack allows administrators to dynamically monitor the health of each accelerator component, collect telemetry data, and perform run-time reconfigurations such as power capping, error logging, and device enablement/disablement. The list of supported devices and platform configurations are as follows:

- ✓ FPGA Device Management
- ✓ GPGPU Device Management
- ✓ Baseboard Device Management
- ✓ NVSwitch Device Management
- ✓ NVLink Device Management
- ✓ PCIe Retimer Device Management
- ✓ PCIe Switch Device Management
- ✓ HMC Platform Support
- ✓ Non-HMC Platform Support

Firmware Updates with Minimal Downtime

Critical firmware updates are notorious for requiring multiple system reboots which can lead to extended system downtime and potential data loss. AMI's solutions extend NVIDIA's parallel firmware update mechanisms to streamline out-of-band firmware deployment. Additionally, enhanced Redfish® support allows real-time system health monitoring via dashboards and HTTP APIs.

Lastly, customers can greatly benefit from complementing this solution with other AMI products such as AMI Data Center Manager (DCM), AMI Tektagon™ (platform security and attestation), or AMI Aptio® V (best-in-class UEFI solution). With AMI, customers are able to save time and effort when beginning or scaling their GPU infrastructure.

NVIDIA Platform Support

- ✓ NVIDIA Grace™
 - NVIDIA MGX™ (C1, C2)
- ✓ NVIDIA Grace™ + NVIDIA Hopper™
 - NVIDIA MGX™ (CG1), GH100, GH200
- ✓ NVIDIA Grace™ + NVIDIA Blackwell™
 - GB200, GB300, GB300 Workstation
- ✓ x86 CPU + NVIDIA GPU
 - NVIDIA HGX™ B300
 - MGX NVL16/32

Interface & Protocol Support

- ✓ Redfish
- ✓ Web-UI
- ✓ MCTP over PCIe
- ✓ SMBPBI over I2C
- ✓ SMBPBI over PCIe
- ✓ PLDM Type 5 FW Update
- ✓ GPGPU AP FW Update
- ✓ FPGA FW Update (via FPGA EROT)
- ✓ FPGA FW Update (via SPI flash)
- ✓ EROT FW Update (GPGPU, FPGA, BOOT)

