

GPU-Accelerated Machine Learning and Inferencing Servers Bring Flexibility and Cost Efficiency to High-Powered Computing



Developed in partnership with USES Integrated Solutions, the AGX Inference Server is an extremely low wattage, high performance AI workstation powered by the NVIDIA Jetson platform. Running NVIDIA's most powerful deep-learning software libraries, this inference server solves the challenge of deploying edge solutions at scale. The on-board 12x Jetson AGX Xavier modules are all connected by a Gigabit Ethernet fabric through a specialized Managed Ethernet Switch with 10G uplink capability.



Includes Out of Band Management (OOBM) via an ARM-based SMARC module to control and monitor each of the Jetson AGX Xavier modules.

FEATURES

- ✓ 12x 11 TFLOPS, 512-core Volta GPU and 64 Tensor Cores
- ✓ 2x 10G SFP+, 2x 1G SFP uplink capability
- ✓ Operating Temperature: 0°C to +50°C (+32°F to +122°F)
- ✓ 1U ATX style redundant power supply
- ✓ Up to 2TB of NVMe storage

SPECIFICATIONS

Processing Modules	<ul style="list-style-type: none"> 12x NVIDIA® Jetson AGX Xavier™ GPU: 512-core Volta GPU with 64 Tensor Cores, 11 TFLOPS (FP16), 22 TOPS (INT8) CPU: 8-core Carmel ARM® v8.2 64-bit CPU, 8MB L2 + 4MB L3 Memory: 32GB 256-Bit LPDDR4x 136.5GB/s Storage: 32GB eMMC 5.1 	Out-of-band Management Module	<ul style="list-style-type: none"> ARM based OOBM Enabling Serial console access, power status monitoring, and power control (ON/OFF) to all 12x AGX Xavier modules OOBM accessible via Ethernet or via its own integrated USB-to-Serial console
Processor Module Carriers	<ul style="list-style-type: none"> Each module carrier will allow up to 4x NVIDIA Jetson AGX Xavier modules to be installed 3x module carriers can be installed in the system for a total of 12 modules 	Internal Embedded Ethernet Switch	<ul style="list-style-type: none"> Vitesse/Microsemi VSC7448 Managed Ethernet Switch Engine (XDG201) CPU: 500 MHz MIPS 24KEc Memory: 4Gb DDR3 SDRAM Storage: 128Mb Serial NOR Flash Multi 1G/10G uplinks, with 12x 1G downstream module connections
Cooling	<ul style="list-style-type: none"> Forced convection via fans on front and rear of system Fans: Delta, 8x 40 x 40 x 28mm, 29 CFM / 2,400 in H2O static pressure 	Internal Array Communication	<ul style="list-style-type: none"> 12x Gigabit Ethernet / 1000BASE-T / IEEE 802.3ab channels All AGX Xavier modules can communicate to all other AGX Xavier modules
External Uplink Connections	<ul style="list-style-type: none"> 2x SFP+ 10G uplink 2x SFP 1G uplinks 	Misc / Additional IO	1x 1GbE OOB management port via RJ-45; 1x USB UART management port; status LEDs
Input Power	100~240 VAC (dual redundant) with 650W output each	Internal Storage	Each AGX Xavier module has its own M.2 NVMe interface (optional)
Operating Temperature	0°C to +50°C (+32°F to +122°F)	Dimensions	Standard 1U rackmount height (1.75 inch / 44.45mm) 25 inch / 635mm Depth