



Housed in a 2U chassis, the Orin NX Inference Server can be outfitted with up to 24x NVIDIA Jetson Orin NX modules.



Developed in partnership with USES Integrated Solutions, the Orin NX Inference Server is an extremely low wattage, high performance deep learning inference server powered by the NVIDIA Jetson Orin NX 16GB module. Within the Server, three processor module carriers house up to eight Jetson Orin NX modules each, and are all connected via a Gigabit Ethernet fabric through a specialized Managed Ethernet Switch developed by Connect Tech with 10G uplink capability (XDG205).



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

## FEATURES

- ✓ 24x 100 TOPS, 1024 GPU CUDA cores with NVIDIA® Ampere™ architecture
- ✓ 4x 10G SFP+ uplink capability
- ✓ Operating Temperature: 0°C to +50°C (+32°F to +122°F)
- ✓ 2U ATX style redundant power supply

## SPECIFICATIONS

<b>Processing Modules</b>	<ul style="list-style-type: none"><li>• 24x NVIDIA® Jetson Orin™ NX</li><li>• GPU: 1024-core CUDA GPU with 32 Tensor Cores, 3.7 TFLOPS (FP16), 100 TOPs (INT8)</li><li>• CPU: 8-core Cortex Arm v8.2 CPU, 2MB L2 + 4MB L3</li><li>• Memory: 16 GB 128-bit LPDDR5   102.4GB/s</li><li>• Storage: Supports External Storage (NVMe) via x4 PCIe</li></ul>	<b>Out-of-band Management Module</b>	<ul style="list-style-type: none"><li>• ARM based OOBM</li><li>• Enabling Serial console access, power status monitoring, and power control (ON/OFF) to all 24x Orin NX modules</li><li>• OOBM accessible via Ethernet or via its own integrated USB-to-Serial console</li></ul>
<b>Processor Module Carriers</b>	<ul style="list-style-type: none"><li>• Each module carrier will allow up to 8x Orin NX modules to be installed</li><li>• 3x module carriers can be installed in the system for a total of 24 modules</li></ul>	<b>Internal Embedded Ethernet Switch</b>	<ul style="list-style-type: none"><li>• Vitesse/Microsemi SparX-5i VSC7558TSN-V/5CC Managed Ethernet Switch Engine (XDG205)</li><li>• CPU: 1 GHz VCore</li><li>• Memory: 8Gb DDR4 SDRAM</li><li>• Storage: 1Gbit Serial NOR Flash</li><li>• Multi 10G uplinks, with 12x 1G downstream</li><li>• Complete TSN feature set (Time-Sensitive Networking)</li></ul>
<b>Cooling</b>	<ul style="list-style-type: none"><li>• Passive heatsinks installed on each Orin NX module</li><li>• Forced convection via fans on front and rear of system</li><li>• Fans: Delta, 5x 70 x 70 x 28 mm, 75 CFM / 2,400 in H2O static pressure</li></ul>	<b>Internal Array Communication</b>	<ul style="list-style-type: none"><li>• 24x Gigabit Ethernet / 1000BASE-T / IEEE 802.3ab channels</li><li>• All Orin NX modules can communicate to all other Orin NX modules</li></ul>
<b>External Uplink Connections</b>	<ul style="list-style-type: none"><li>• 4x SFP+ 10G uplink</li></ul>	<b>Misc / Additional IO</b>	1x 1GbE OOB management port via RJ-45; 1x USB UART management port; status LEDs
<b>Input Power</b>	100~240 VAC (dual redundant) with 1000W output each	<b>Internal Storage</b>	Each Orin NX module has its own M.2 NVMe interface
<b>Operating Temperature</b>	0°C to +50°C (+32°F to +122°F)	<b>Dimensions</b>	Standard 2U rackmount height (3.5 inch / 88.9mm) 25 inch / 635mm depth