

SKY-TESL-A100-80P


SKY-TESL-A800-80P

SKY-TESL-A30-24P

NVIDIA® Tesla® A100 80GB

NVIDIA® Tesla® A800 80GB

NVIDIA® Tesla® A30



Features

- NVIDIA Ampere GPU architecture
- Compute-optimized GPU
- 6912 / 6912 / 3584 NVIDIA® CUDA® Cores
- 432 / 432 / 224 NVIDIA® Tensor Cores
- 80GB HBM2e / 24GB HBM2 memory with ECC
- Up to 1,935 / 1,935 / 933 GB/s memory bandwidth
- Max. power consumption: 300W / 300W / 165W
- Graphics bus: PCI-E 4.0 x16
- Thermal solution: Passive

Introduction

NVIDIA® Tesla® A100 80GB (SKY-TESL-A100-80P), Tesla A800 80GB (SKY-TESL-A800-80P) and Tesla A30 (SKY-TESL-A30-24P) PCIe cards are compute-optimized GPUs built on the NVIDIA Ampere architecture with dual-slot 10.5-inch PCI Express Gen4 interface in a passive heatsink cooling design suitable for data centers. Combining NVIDIA Gen3 tensor cores and HBM2e/HBM2 memory, they provide a high-performance computing solution. Supporting a multi-instance GPU (MIG) feature, which guarantees quality of service (QoS) with secure, partitioned hardware, they allow maximum utilization of GPU resources. The NVIDIA NGC™ catalog provides software, libraries, and optimized AI models and applications to complete data center solutions. With cutting-edge features and technologies, NVIDIA Tesla A100 80GB, A800 80GB and A30 are perfect for AI deep learning training and inference, data analytics, and high-performance computing (HPC) applications. NVIDIA Tesla is the first choice for high-standard computing solutions in enterprise and science deployments.

Specifications

Product Name	Tesla A100 80GB	Tesla A800 80GB	Tesla A30
Part Number	SKY-TESL-A100-80P	SKY-TESL-A800-80P	SKY-TESL-A30-24P
GPU Architecture	Ampere	Ampere	Ampere
GPU Memory	80GB HBM2e	80GB HBM2e	24GB HBM2
Memory Bandwidth	1,935 GB/s	1,935 GB/s	933GB/s
NVIDIA CUDA Cores	6912	6912	3584
Tensor Cores	432	432	224
Single-Precision Performance	19.5 TFLOPS	19.5 TFLOPS	10.3 TFLOPS
Double-Precision Performance	9.7 TFLOPS	9.7 TFLOPS	5.2 TFLOPS
Fast FP64	Yes	Yes	Yes
System Interface	PCI Express 4.0x16	PCI Express 4.0x16	PCI Express 4.0x16
Max Power Consumption	300W	300W	165W
Power Connector	8-Pin CPU	8-Pin CPU	8-Pin CPU
Thermal Solution	Passive	Passive	Passive
Multi-Instance GPU	Up to 7	Up to 7	Up to 4
Form Factor	4.4 inches H x 10.5 inches L dual slot, full height	4.4 inches H x 10.5 inches L dual slot, full height	4.4 inches H x 10.5 inches L dual slot, full height
NVLink Support	3 NVLINK Bridges for 2 GPUs, 600GB/s	3 NVLINK Bridges for 2 GPUs, 400GB/s	1 NVLINK Bridge for 2 GPUs, 200GB/s
Media Acceleration	1 JPEG Decoder, 5 Video Decoder	1 JPEG Decoder, 5 Video Decoder	1 JPEG Decoder, 4 Video Decoder
Display Connectors	Headless Design	Headless Design	Headless Design