

# NVIDIA® L40S

## SPECIFICATIONS<sub>1</sub>

<b>GPU Architecture</b>	NVIDIA Ada Lovelace Architecture
<b>GPU Memory</b>	48GB GDDR6 with ECC
<b>Memory Bandwidth</b>	PCIe Gen4x 16: 64GB/s bidirectional
<b>CUDA™ Cores</b>	18176
<b>RT Cores</b>	142
<b>Tensor Cores</b>	568
<b>RT Core Performance</b>	212 TFLOPS
<b>FP32</b>	91.6 TFLOPS
<b>TF32 Tensor Core</b>	366 TFLOPS
<b>BFLOAT16 Tensor Core</b>	366   733 <sub>2</sub> TFLOPS
<b>FP16 Tensor Core</b>	366   733 <sub>2</sub> TFLOPS
<b>FP8 Tensor Core</b>	733   1466 <sub>2</sub> TFLOPS
<b>Peak INT8 Tensor</b>	733   1466 <sub>2</sub> TOPS
<b>Peak INT4 Tensor</b>	733   1466 <sub>2</sub> TOPS
<b>Form Factor</b>	4.4" (H) x 10.5" (L), dual slot
<b>Display Ports</b>	4 x DisplayPort 1.4a
<b>Max Power Consumption</b>	350W
<b>Power Connector</b>	16-pin
<b>Thermal</b>	Passive
<b>Virtual GPU (vGPU) Software Support</b>	Yes
<b>NVENC   NVDEC</b>	3x   3x (Includes AV1 Encode and Decode)
<b>Secure Boot with Root of Trust</b>	Yes
<b>NEBS Ready</b>	Level 3
<b>MIG Support</b>	No

1. Preliminary specifications, subject to change.

2. Specifications with sparsity.