



# NVIDIA L40

Delivering unprecedented visual computing performance for the data center.



From virtual workstation application to large-scale modeling and simulation, modern visual computing and scientific workflows are growing in both complexity and quantity. Enterprises need data center technology that can deliver extreme performance and scale with versatile capabilities to conquer the diverse computing demands of these increasingly complex workloads.

The NVIDIA® L40 GPU delivers unprecedented visual computing performance for the data center, providing next-generation graphics, compute, and AI capabilities. Built on the revolutionary NVIDIA Ada Lovelace architecture, the NVIDIA L40 harnesses the power of the latest generation RT, Tensor, and CUDA cores to deliver groundbreaking visualization and compute performance for the most demanding data center workloads.

## Powered by the NVIDIA Ada Lovelace Architecture

### Third-Generation RT Cores

Enhanced throughput and concurrent ray-tracing and shading capabilities improve ray-tracing performance, accelerating renders for product design and architecture, engineering, and construction workflows. See lifelike designs in action with hardware-accelerated motion blur to deliver stunning real-time animations.

### Fourth-Generation Tensor Cores

Hardware support for structural sparsity and optimized TF32 format provides out-of-the-box performance gains for faster AI and data science model training. Accelerate AI-enhanced graphics capabilities, including **DLSS**, delivering upscaled resolution with better performance in select applications.

### Large GPU Memory

Tackle memory-intensive applications and workloads like data science, simulation, 3D modeling, and rendering with 48GB of ultra-fast GDDR6 memory. Allocate memory to multiple users with vGPU software to distribute large workloads among creative, data science, and design teams.

### Data-Center Ready

Designed for 24x7 enterprise data center operations with power-efficient hardware and components, the NVIDIA L40 is optimized to deploy at scale and deliver maximum performance for a diverse range of data center workloads. The L40 includes secure boot with root of trust technology providing an additional layer of security, and is NEBS Level 3 compliant to meet the latest data center standards. Packaged in a dual-slot, passively cooled and power-efficient design, the L40 is available in a wide variety of NVIDIA-Certified Systems™ from leading OEM vendors.

## Accelerate Next-Generation Workloads

- > NVIDIA Omniverse™ Enterprise
- > Rendering and 3D Graphics
- > High-Performance Virtual Workstations with NVIDIA RTX™ Virtual Workstation (RTX vWS) Software
- > AI Training and Data Science
- > Streaming and Video Content

## Technical Specifications\*

<b>GPU Architecture</b>	NVIDIA Ada Lovelace architecture
<b>GPU Memory</b>	48GB GDDR6 with ECC
<b>Memory Bandwidth</b>	864GB/s
<b>Interconnect Interface</b>	PCIe Gen4x16: 64GB/s bi-directional
<b>NVIDIA Ada Lovelace architecture-based CUDA Cores</b>	18,176
<b>NVIDIA third-generation RT Cores</b>	142
<b>NVIDIA fourth-generation Tensor Cores</b>	568
<b>RT Core performance TFLOPS</b>	209
<b>FP32 TFLOPS</b>	90.5
<b>TF32 Tensor Core TFLOPS</b>	90.5   181**
<b>BFLOAT16 Tensor Core TFLOPS</b>	181.05   362.1**
<b>FP16 Tensor Core</b>	181.05   362.1**
<b>FP8 Tensor Core</b>	362   724**
<b>Peak INT8 Tensor TOPS</b>	362   724**
<b>Peak INT4 Tensor TOPS</b>	724   1448**
<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>	300W
<b>Power Connector</b>	16-pin
<b>Thermal</b>	Passive
<b>Virtual GPU (vGPU) software support</b>	Yes
<b>vGPU Profiles Supported</b>	See <a href="#">Virtual GPU Licensing Guide</a> <sup>1</sup>
<b>NVENC   NVDEC</b>	3x   3x (Includes AV1 Encode & Decode)
<b>Secure Boot with Root of Trust</b>	Yes
<b>NEBS Ready</b>	Level 3
<b>MIG Support</b>	No
<b>NVLink Support</b>	No

\* Preliminary specifications, subject to change.

\*\* With Sparsity.

## Ready to get started?

To learn more about the NVIDIA L40 GPU, visit:

[www.nvidia.com/L40](http://www.nvidia.com/L40)

1. Coming in a future release of NVIDIA vGPU software.

© 2023 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA-Certified Systems, NVIDIA Omniverse and NVIDIA RTX are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 2436245. JAN23

