

# NVIDIA LICENSING ON HC3 WITH GPU

## OVERVIEW AND DISCLAIMER

Answers to general questions around NVIDIA licensing and how it relates to common HC3 implementations with GPU. This FAQ provides helpful terminology and a high-level overview of information from the NVIDIA licensing guide, as well as other sites, to help you get started on the right track with NVIDIA licensing.

Scale Computing is in no way responsible for any issues or problems arising from the information found in this FAQ. Contact NVIDIA for the most current information or any detailed licensing questions or concerns. This document is in no way authoritative on NVIDIA licensing requirements, nor will it provide indemnity in the event NVIDIA conducts a licensing or entitlement audit or review.

## FREQUENTLY ASKED QUESTIONS

**Q: What NVIDIA GPUs does Scale Computing support?**

**A:** HC3 supports the NVIDIA Tesla T4 in 2 and 5 GPU configurations as of the publish date of this guide. For the latest offerings, please see the [product family details](#). Details of the Tesla T4 are below.

	NVIDIA T4
Recommended Use case	Quadro vDWS - Performance Optimized (Entry to mid); GRID vPC - Density Optimized; vCS – Compute Optimized
Number of GPUs	1 NVIDIA Turing™ TU104
Total CUDA Cores	2,560
Tensor Cores	320
RT Cores	40
Total Memory Size	16 GB GDDR6
Max GPU Power	70 W
Form Factor	PCIe 3.0 Single Slot
Board Dimensions	2.7"×6.6"
Cooling Solution	Passive

**Q: Do I need to license NVIDIA for use with HC3 GPU nodes? What NVIDIA licensing is included with an HC3 node or cluster purchase with GPU?**

**A:** Yes. Users will need to purchase licensing from NVIDIA in order to utilize virtual GPUs (vGPUs) on HC3 nodes with GPUs. The purchase of a Scale Computing HC3 node or cluster does not include any licensing outside of the HyperCore operating system (HC3 Software). Microsoft, Oracle, NVIDIA, and other licenses are not included or resold through Scale Computing.



**Q: What NVIDIA licensing options or models are available?**

**A:** Licensing for NVIDIA depends highly on the use case for which it is being deployed. The following licensing options are available:

- **NVIDIA GRID Virtual Applications (GRID vApps)** - For organizations deploying RDSH or app streaming or session-based solutions. Designed for PC-level applications and server-based desktops.
- **NVIDIA GRID Virtual PC (GRID vPC)** - For users who want a virtual desktop but need great user experience leveraging PC Windows applications, browsers, and high definition video.
- **NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS)** - For users who want to be able to use remote professional graphics applications with full performance on any device, anywhere.
- **NVIDIA Virtual Compute Server (vCS)** - For compute-intensive server workloads, such as artificial intelligence (AI), deep learning, or high-performance computing (HPC).

**Q: How much are the various licensing options?**

**A:** NVIDIA provides both subscription and perpetual license options in various price points. USD examples are provided below.

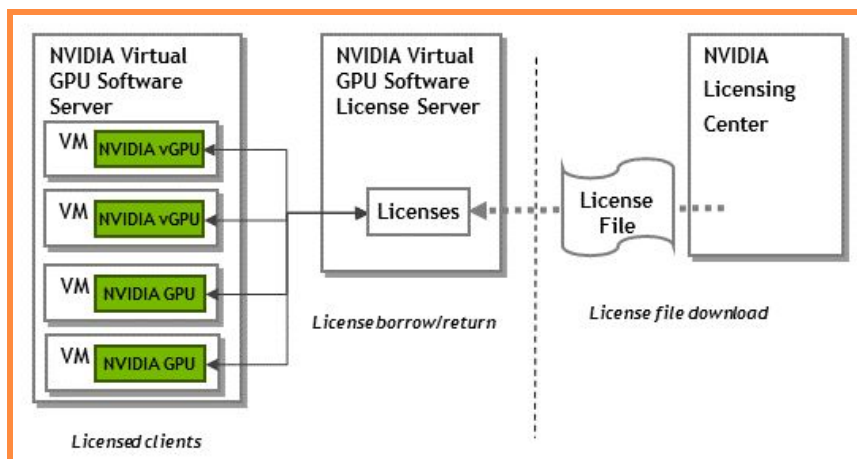
ANNUAL SUBSCRIPTION		PERPETUAL LICENSE	
			
GRID Virtual Applications	\$10 per concurrent user subscription	GRID Virtual Applications	\$20 perpetual license \$5 SUMS per year
GRID Virtual PC	\$50 per concurrent user subscription	GRID Virtual PC	\$100 per CCU \$25 SUMS per year
Quadro Virtual Data Center Workstation	\$250 per concurrent user subscription	Quadro Virtual Data Center Workstation	\$450 perpetual license \$100 SUMS per year
NVIDIA Virtual Compute Server	\$150 per GPU subscription	NVIDIA Virtual Compute Server	Perpetual License not available

Annual subscription includes software license and SUMS.

Perpetual License includes indefinite software license; SUMS is required and is available in four or five year increments. One year SUMS available only for renewals.

**Q: How are licenses applied to the VMs on HC3?**

**A:** Licensed vGPU functionalities are activated during guest OS boot by the acquisition of a software license served over the network from an NVIDIA vGPU software license server. The license is returned to the license server when the guest OS shuts down. The diagram shows how NVIDIA vGPU software license files are downloaded from the NVIDIA Licensing Portal to the license server, and how licensed clients borrow licenses from the server.



**Q: Are there requirements for where the licensing server can be deployed?**

**A:** The licensing server can be deployed as a physical or virtual server. The licensing server does not need to be hosted on the system with the GPUs or any virtual guests which require the licensing. The only requirement is to ensure the vGPU guests have network access to the licensing server.

**Q: Can I set up more than one licensing server to provide redundancy?**

**A:** Yes. The licensing server and guest drivers have provisions for redundant server configurations. See the [NVIDIA Virtual GPU Licensing Server User Guide](#) for setup details.

## OTHER RESOURCES

- [NVIDIA Licensing Overview](#)
- [NVIDIA Virtual GPU Software Packaging, Pricing, and Licensing Guide](#)
- [Virtual GPU Software License Server User Guide](#)