

V-Ray GPU Frequently Asked Questions

Which features are supported on the GPU?

The code base for V-Ray GPU differs from the code base for the CPU engine. The GPU code base is optimized for GPU hardware and supports a subset of the features of CPU rendering. For a complete list of supported features, please see the [V-Ray GPU Supported Features](#) page.

Can I choose which devices to be used for GPU rendering?

Yes. There are instructions how to do that in the [GPU Setup](#) section.

How can I check which devices on my system can be used for GPU rendering?

Why isn't my GPU device(s) listed?

Check the list of devices in the GPU Devices list accessed via the button next to the GPU Engine button in the Asset Editor's Settings tab. If you don't see your device on the list there (the list is empty), please check whether it is [supported](#).

Why does the UI of my OS becomes sluggish when I use GPU?

- If you have multiple GPUs, you can speed up screen refresh time by removing the GPU used for monitor/viewport redraw from the devices that GPU uses for rendering.
- If you have only one GPU on your system, you could try reducing the value for **Max Trace Depth** (accessible from Asset Editor > Settings > Render Parameters > Optimizations). This will break up the data passed to the GPU into smaller chunks so that user interface requests can be processed faster. However, this might reduce rendering speed.

V-Ray will display a warning if the GPU used for monitoring is participating in the rendering process.

How do I reduce the network bandwidth when using GPU with a heavy DR setup?

If you use many GPUs in DR nodes you may find that the network can become a bottleneck at some point, since the GPUs are producing data very fast. Increasing "Max Trace Depth" / "Max Ray Intensity" increases the size of the workload chunks that are being given to the DR nodes and helps reduce the communication between the client and the servers (you can try something like 192/32 or 256/64).

How do I check the memory usage on the GPU, and reduce the GPU memory needed for my scene?

You can use free 3rd party tools like [MSI Afterburner](#) and [EVGA Precision](#) to monitor both GPU memory usage and utilization. V-Ray GPU also reports how much memory it uses for textures/geometry/light cache/etc in the V-Ray Progress log accessible via a button at the bottom-right corner of the VFB.

Can GPU be used as production render?

Yes, but within its [limitations](#).

How do I setup V-Ray GPU as my production renderer?

What are the differences between Interactive and Production mode?

You can set V-Ray GPU as production renderer by selecting the **GPU** engine from the **Render** rollout of the **Asset Editor's Settings** tab and disabling the **Interactive** toggle.

The main differences between Interactive and Production rendering is that Production allows you to render animations, and automatically save the render output after the process has finished, while Interactive rendering does not. Once the final noise threshold is reached in Interactive rendering, the process will remain active awaiting further changes to the scene. Manually stopping the render process does not produce automatically an image written to disk. Another difference is that in Production mode V-Ray will calculate and use Light Cache for GPU GI calculations if it has been set as secondary GI engine. In Interactive mode V-Ray GPU always uses Brute Force for both primary and secondary GI engine, unless Light Cache is used from file.

How do I use Light Cache with GPU?

Light Cache can be used as secondary GI engine with V-Ray GPU for production rendering. Brute Force is always used as the primary GI engine. In Interactive mode, however, Brute Force is always used for primary and secondary GI engine, regardless of the GI engine settings, unless you load your Light Cache from a file. The reason for that is that the Light Cache is view dependent, and recalculating it every time the camera changes its position in Interactive mode decreases performance.

Why do some bump maps look different with GPU?

Procedural textures are supported for GPU bump maps, but because V-Ray GPU is a very different engine, the bump map itself can look a bit different compared to V-Ray.

Larger scenes cause a "Failed to allocate memory for buffer" error when using the GPU engine. Why?

If you get the following error message, this means that some or all of your video cards do not have sufficient memory to load the scene. To solve this problem, you can enable texture resizing for GPU rendering and set a texture size limit. This will cause V-Ray GPU to use smaller textures for rendering to conserve memory. See more information on GPU texture resizing.

Why is the GPU rendering not much faster than the CPU rendering?

The difference in render speeds depends on the video cards and the CPUs that are compared, as well as on the scenes used as benchmarks. It is normal to achieve a big speed boost with the GPUs compared to the CPU with some scenes and to have equal performance with others.

Can I use multiple systems with GPU devices to speed up my rendering?

Yes, you can speed up your rendering using Distributed Rendering (DR) on multiple systems with CUDA enabled devices. This can be easily set up using V-Ray Swarm.

Keep in mind that some drivers may have limitations and may not be able to use GPUs if there is no monitor attached to the machine or you are logged in via Remote Desktop (for example). Check the GPU vendor documentation for more information.

How to start Interactive rendering - with or without the VFB?

There are two ways to start interactive rendering with V-Ray:

1. From the Toolbar or Asset Editor **Render Interactive** button;
2. From the Toolbar **Viewport Render** button.

Can I use V-Ray GPU on Mac OS? How?

V-Ray GPU is not officially supported on macOS.

It works only with C++/CPU devices. V-Ray GPU can still be used in distributed rendering where a macOS machine runs the CUDA engine on a CPU device together with Windows/Linux machine(s) running CUDA engine on GPU device(s).