Getting Started With Arnold GPU

Supported Features

Arnold GPU supports a set number of Arnold features, including arbitrary shading networks, SSS, hair, atmospherics, instancing, and procedurals. See here for a detailed list of Arnold GPU features and known limitations.

System Requirements

Arnold GPU works on NVIDIA GPUs of the Ampere, Turing, Volta, Pascal, and Maxwell architectures. Multiple GPUs will improve performance, and NVLink can be used to connect multiple GPUs of the same architecture to share memory (On Windows, we recommend enabling SLI as well).

See the full list of supported GPUs.

Recommended NVIDIA drivers:

- **Linux** 470.74 or higher
- **Windows** GeForce: 472.12 or higher, Quadro: 472.12 or higher
- macOS is not supported

NVIDIA Driver Downloads

Pre-populating the GPU cache

The very first time you render with the GPU, the GPU renderer has to create a cache for the types of shaders and objects used. This can delay the time to the first pixel for your first render. To avoid these one-time delays, we recommend doing some of this in advance by pre-populating the cache before you do any renders as this will fill in the cache with a subset of the most common shader/object combinations. The Arnold plugins (like MtoA) have *Pre-populate GPU Cache* menu command, and kick has a -gpu_warm flag. Note that pre-populating the cache can take up to 15 minutes. The cache only needs to be re-populated after installing a new Arnold version, updating to a new NVIDIA driver, or changing the hardware configuration of GPUs on the system.

The GPU cache is based on the Arnold core version, NVIDIA driver version, and GPU(s) the software is executed on. So if one of them changes then you will need to run the *Pre-populate Cache* again, to avoid (one-time) delays during rendering.

Selecting a Render Device

You can easily switch between CPU and GPU with a single click. You'll find a new Render Device setting in the Render Settings > System section in the Arnold plugins.

You can use `CUDA_VISIBLE_DEVICES` (environment variable) to limit which GPUs Arnold (and any other Cuda-based application) can see.

Matching Noise on CPU and GPU

Matching noise can take a little experimentation because **Arnold GPU uses *Camera (AA)* sampling only**. We recommend you also use *Adaptive* sampling. Here are some guidelines:

- Set the *Max. Camera (AA)* in the range of 30 to 50 (depending on the scene, you might go closer to 100). In general, the max samples should be a large value. A large max samples means that the quality is controlled by the noise falling under the threshold, instead of by clamping to the max AA.
- Set the *Adaptive Threshold* to something like 0.015 or 0.02. For a noise-free render, lower the threshold value, maybe even as far as 0.010.
- Set the *Camera (AA)* samples to around 3 or 4. One of the few reasons to go higher with AA is for motion blur. The higher the number of *Camera (AA)* samples, the less of a speedup you'll get from adaptive sampling.