

Distributed Rendering

This page provides information on Distributed Rendering with Chaos Phoenix.

Overview

Distributed Rendering (DR) is a technique for dispersing rendering jobs within a single frame across multiple computers in a network. There are different approaches to doing this but the main concept is to reduce the render times by dividing different parts of the rendering task and giving each participant different parts to complete. The most common way to do this is to divide the frame to be rendered into small regions and giving each rendering machine one after another until the entire frame or job is complete. Then the results are combined into the final image.

Note that distributed simulations in which different parts of the same simulation are calculated simultaneously across several machines are currently **not implemented**. You can use the Simulation licenses to run simulations on separate machines without accessing the user interface and without engaging a GUI license. For more information on how to submit simulations that use only a Simulation license and how to do advanced scripting for more complicated tasks, see [Simulation Without a GUI License](#).

Setting up Phoenix Distributed Rendering using V-Ray

For more information on DR using V-Ray, see the [Set up Distributed Rendering](#) page.

Troubleshooting

If you're able to render locally but when using DR machines, the Phoenix fluid is missing from the buckets, please try the following:

- Try using the latest versions of Phoenix and V-Ray.
- Confirm that the version numbers are the same on all DR machines.
- Make sure the simulation caches are shared on a network location that all the machines can see. Note: They all must have permissions to access this network folder. *For more details, see [Distributed Rendering of Fluid Caches for Phoenix](#) below.*
- Confirm that your V-Ray Spawner is started as an administrator on the remote machines. *For more details, see [Mapped Drives and UNC Paths for Texture Maps and Other Rendering Assets](#) below.*
- Make sure your V-Ray Spawner is run as the same account as the one you are testing the permissions with. Your local account might be able to see the shared network folder but the spawner might be started with a different user account which doesn't have proper permissions.

Distributed Rendering of Fluid Caches for Phoenix

A common problem when setting up Distributed Rendering is that by default Phoenix looks for the cache files in the same directory as the scene file using the **\$(scene_path)** keyword. However, at the start of a network render, the scene file is copied to all render machines on the network to a new location (e.g. C:\Users\user\AppData\Local\backburner\Jobs\). While the cache files are not automatically sent to the host machine(s). They are not sent because they may be really huge and may overload the network and the disk space of the host, and would take too much time to transfer, which will greatly reduce the performance of Distributed Rendering. This is why when rendering begins, if the hosts are looking for the cache files in the same directory where the scene file is, the cache files won't be found.

The solution is to move the cache sequence to a shared folder on the network or a mapped network drive and set its path in the **Input** panel, using a network-visible **UNC input path** (a path that starts with \\). You can also browse from the **Input** panel's path options.

If you are not using a UNC path or a mapped network drive, and instead your path points to a drive on your local computer or is one of the Phoenix keywords (**\$(scene_path)**, **\$(implicit)**, **\$(scene_dir)**, etc...), you will receive a message that states "You are using local machine Input path with distributed rendering!".

Mapped Drives and UNC Paths for Texture Maps and Other Rendering Assets

If you are using mapped drives or UNC paths for your scene assets, you should set up a "Log On" account for the V-Ray spawner service, otherwise, the textures or any other external assets won't be accessible. To set up the Log On account, open the **Services** applet from the Control panel (*Control panel > Administrative Tools > Services*) and find "VRaySpawner 20nn" service, then right-click to expand the drop-down menu and choose **Properties**. Go to the **Log On** tab and instead of the default *Log on as local system account* choose **This account** and put the name of an account that has adequate privileges on the machine and access to the shared resources over the network.

Alternatively, you can use the automatic asset transfer mechanism of V-Ray as described above to send the required assets from the client machine to the render servers.

