# Setup V-Ray for Rhino to use the Render Farm:

V-Ray for Rhino provides a distributed network rendering capability for use within a V-Ray render session in Rhino. It is initiated through the standard V-Ray Options dialog box within Rhino, and will distribute the on-screen render job via buckets to the various render nodes selected.

Please note that this is a NOT an offline render queue, but rather a network assist for interactive on-screen rendering. V-Ray for Rhino does not offer an offline render queue. You must remain on the local for the render to continue working. You cannot close down or log off.

To setup access from a SARC workstation or your laptop computer, you need only have Rhino and V-Ray for Rhino installed. You also must be on the UVA network, either directly or through a VPN (UVA Anywhere) connection.

## Setup the Nodes

To render using the Render Farm grid, you need to do two things to setup:

1. **Turn on Distributed Rendering**
   Within the standard V-Ray Option editor window (where you setup most V-Ray options), find the **System** block.

   Expand it, and about half-way down this option block, you will see a checkbox to turn Distributed Rendering On. Check this box.

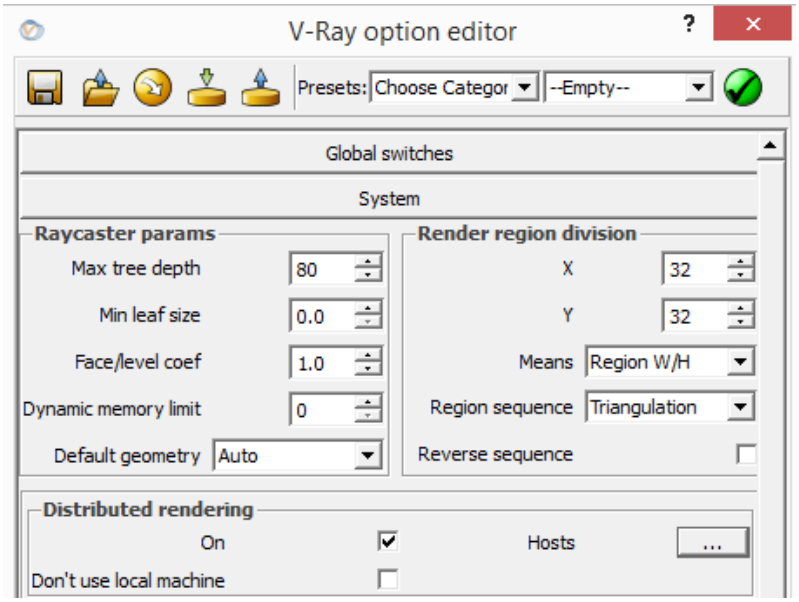


2. **Connect to the Render Farm Nodes**
   Now click the **Hosts …** button, and V-Ray will bring up a dialog box in which to add the render node addresses to connect to.

Use the **Add Server** button to add the IP Address of each node.

The node list is:

renderfarm01.arch.virginia.edu
renderfarm02.arch.virginia.edu
… through …
renderfarm08.arch.virginia.edu

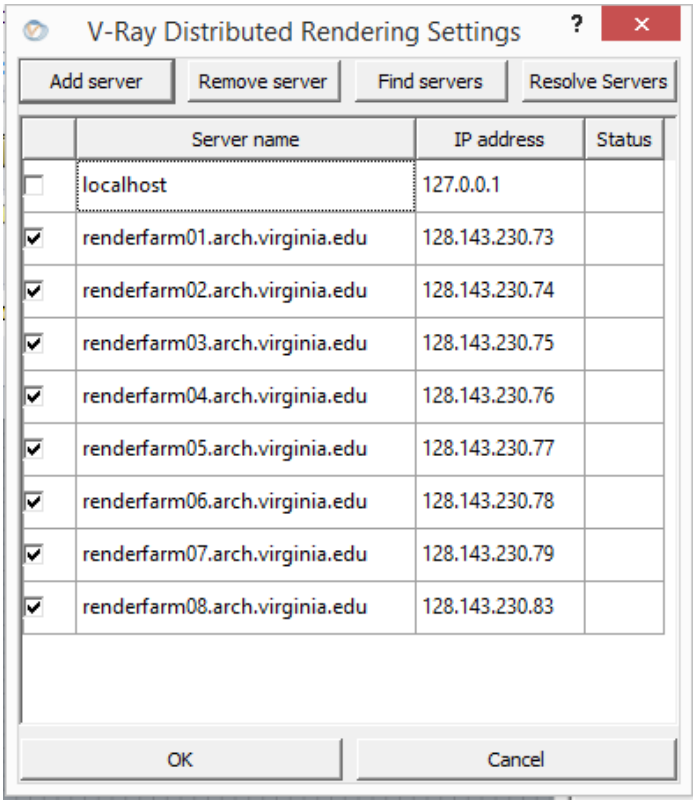


Once done, click the checkbox next to each node, then click the **Resolve Servers** button. It will find the servers and return the IP addresses (which is the same as what you put in, but this verifies that it is resolving them).

Note that this will only work from within the UVA network or VPN session. These IP addresses are not visible to the public internet.

Click **OK**, then close the Options dialog box.

## Render

Now back in the usual Rhino / V-Ray rendering workflow, click the Render icon as you normally would.

The Render Status window will pop up as it normally does, but this time it will include information about connecting to the render hosts that you just specified.

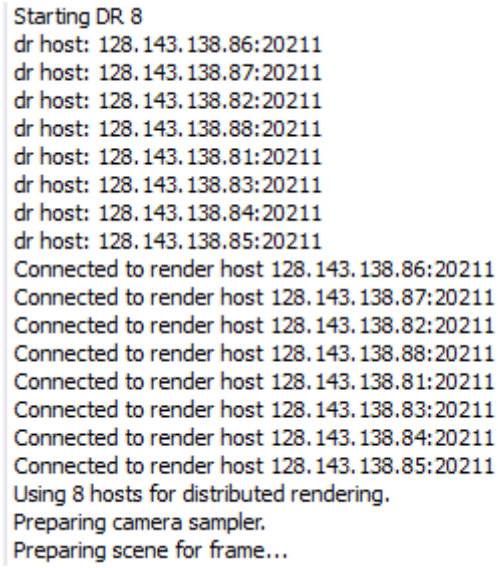


The render job will begin, and you will see buckets with render node names assigned to them. The local computer will be included in this list also, and it will likely be the only assigned node you see for a short time. Depending on the size and complexity of the model and materials, it

may take a few minutes for Rhino / V-Ray to send the geometry and materials to each of the render nodes. The time that this takes can vary widely, based on the file being sent and the network traffic between you and the server. If your job takes too long to send, you may find that the rendering will be completed entirely on the local machine anyway. Again, this depends on the geometry.

Once the job does get to the nodes, you will see their names start to appear on the buckets of the rendering. You should now see a significant speed up in the progress of the rendering once this happens.
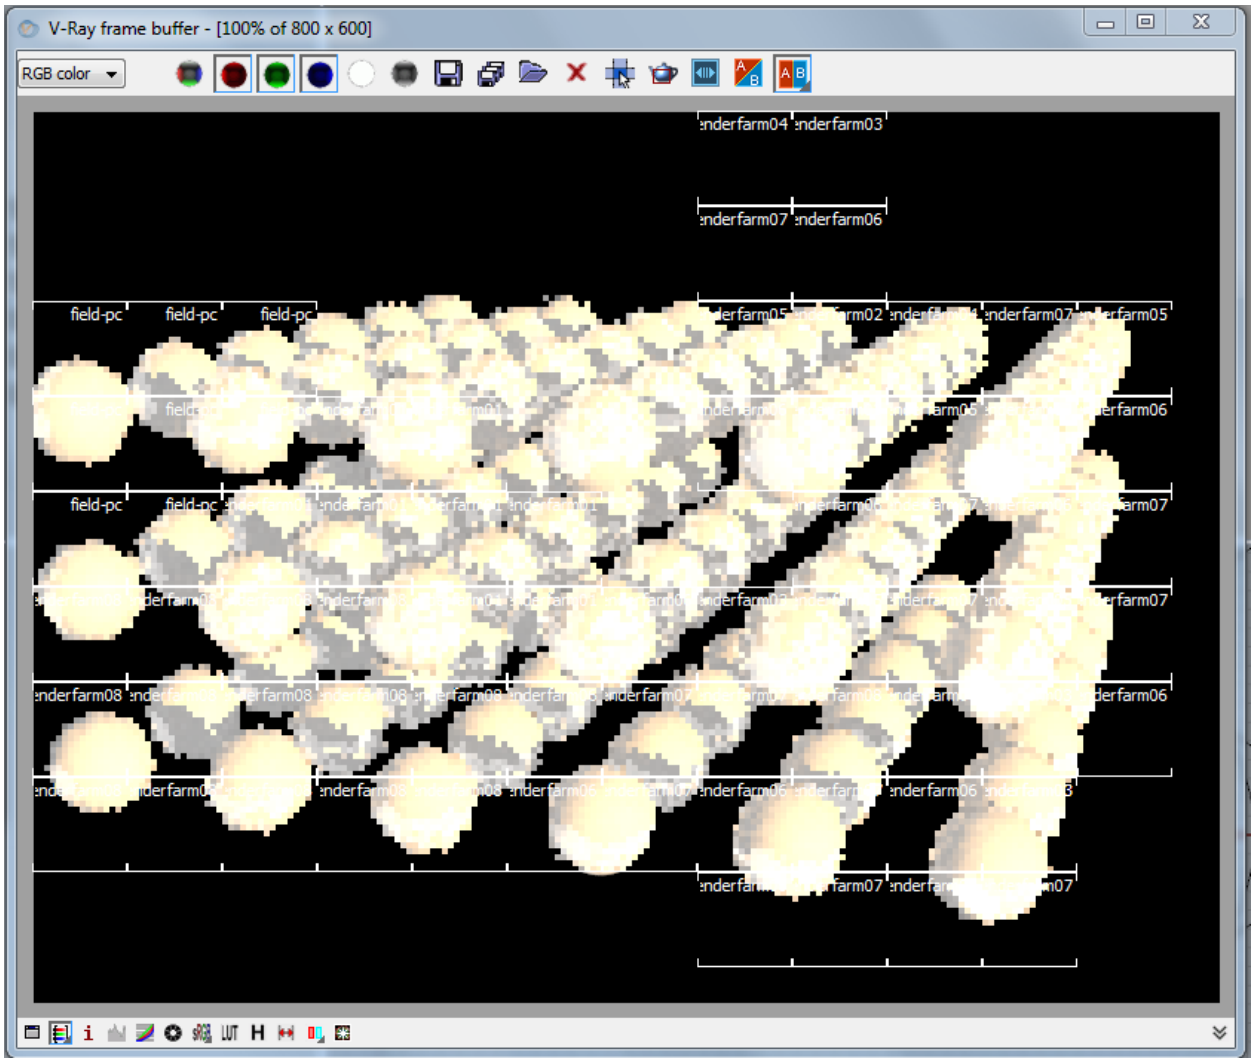
Once the rendering is complete (all buckets have rendered), save the render image as you normally would. There is no special procedure for using the network rendering at this stage. Distributed Network Rendering is simply a real-time assist to the usual V-Ray for Rhino rendering workflow.