

# FireSmoke Simulation

This page provides information on the Simulation rollout for a FireSmokeSim object.

## Overview

The Simulation rollout contains the main panel for controlling the simulation process. It also displays statistical information like simulation times, and the content of the cache file for the current frame. The content can be grid channels such as Temperature, Velocity, etc., as well as particle groups. You can select which channels to export to the cache files from the [Output rollout](#), while the particle groups are automatically created by the Simulator.

Many simulations require a long time to calculate and it's very convenient to let them run during the night. However, you still have to render the result in the morning, and this also consumes a lot of time. In this case, the [Phoenix Scripting](#) system can be helpful, since it enables you to execute any action at the end of the simulation, including rendering. You simply need to enable "Use Script" from the Simulation rollout, and open the script text. For more information, see the Nightly Simulation and Rendering section on the [Tips and Tricks](#) page.

UI Path: `[[Select Fire Smoke Simulator]] > Modify panel > Simulation rollout`

## Actions

**Expand** – Opens a floating dialog that contains the selected rollout and automatically folds the command panel rollout.

**Re-Center** - Resets the position of the floating rollout.

**?** – Opens up the help documents for the Fire/Smoke Simulation options.

**Start, Pause, Stop, Resume** | [A\\_StartSim](#), [A\\_StopSim](#) – Starts, pauses and stops the simulation. Once the simulation has been paused, the option to **Resume** becomes available. Note that you can stop a running simulation using the **Shift+Escape** key combination, or with the **Escape** key, if the option to do so is enabled in the [PhoenixFD Global Preferences](#).

**Load & Start** | [A\\_StartSim](#) – Loads the data from a single cache as an initial state for the simulation and starts simulating from the **Start Frame** the same way the **Start** action works. The cache can be of the *.aur*, *.vdb*, or *.f3d* formats and so can be imported from a 3rd party software into Phoenix. The loaded grid will be resized to fit the grid's dimensions at the **Start Frame**. The loaded cache does not need to be a **Backup** frame or to contain velocity at all - the simulator will load any available channels from the cache file and the rest of the simulated channels will be empty as when starting a brand new simulation. If you wish to smoothly continue a simulated sequence from a given frame, please use the **Restore** command instead.

**Restore** | [A\\_StartSim](#) – Continues a simulation from the currently viewed frame on the timeline, or the latest **Backup** frame containing the full simulation state up to the current timeline frame, which is controlled by the [Backup Interval](#) parameter in the [Output rollout](#). This way, a Phoenix simulation that has been stopped previously can be resumed from the point you left off, even after the software has been closed and reopened.

The ability to restore a sim is only possible when the full internal state of a Phoenix simulation has been exported. This is controlled by the Backup Interval parameter in the Output rollout, which determines how often the full simulation state is backed up. Frames with the full state are called **Backup frames**.

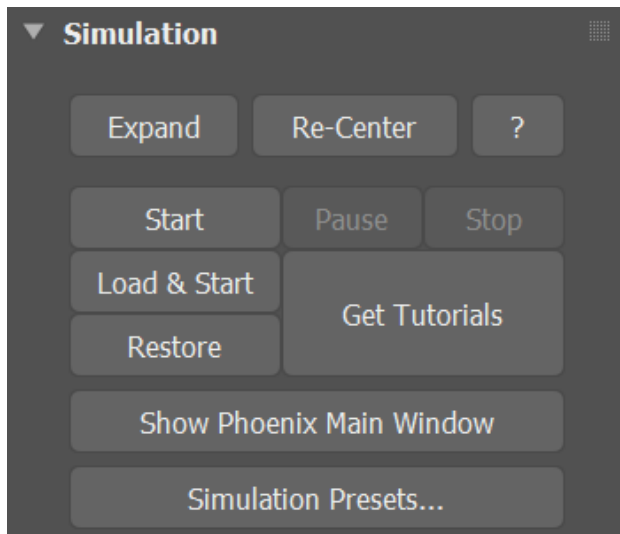
When a **Backup** frame is currently viewed, the text "Can Restore From Here" appears in the Cache File Content list. The **Restore** command finds the latest **Backup** frame up to the current timeline frame, and continues the simulation from there. Note that restoring does not work for cache sequences imported from 3rd party software.

**Get Tutorials** – Opens the [Tutorials and Examples](#) page in the web browser.

**Phoenix Help** – Opens up the help documents for the [Fire/Smoke Simulation](#).

**Show Phoenix Main Window** - Opens a floating dialog where floating rollouts of Simulators can be docked. The layout of this window gets saved to the .MAX scene files.

**Simulation Presets...** | [A\\_LoadRenderPreset](#) – Displays a drop-down menu to either load or save *.tpr* or *.slog* files, which hold settings for this simulation.



## Parameters

**Start Frame** | *startframe* – Explicitly sets the Start frame of the Simulation. This can also be a negative number. Only available when **Timeline** (Start Frame) is disabled.

**Timeline** | *start\_from\_timeline* - When enabled, the Simulation will run from the Timeline Start Frame.

**Stop Frame** | *stopframe* – Explicitly sets the End frame of the Simulation. This can also be a negative number. Only available when **Timeline** (Stop Frame) is disabled.

**Timeline** | *stop\_from\_timeline* - When enabled, the Simulation will run to the end of the Timeline.

**Use GPU to Aid Simulation** | *useGPU* – When enabled, runs suitable phases of the simulation on the GPU instead of the CPU, which can help the simulation speed. Currently only the PCG Fluidity (Conservation) method of Fire/Smoke simulations takes advantage of this.

The PCG fluidity on CPU and the PCG fluidity on GPU will not produce identical results - the fluid will slowly start moving in slightly different directions between the two implementations.

If the GPU can't be initialized, Phoenix will automatically fall back to CPU and will indicate this in the Phoenix log file. The reason the GPU is not initialized may be that there is no GPU, or the GPU is old, or the GPU drivers need to be updated.

**Render On Each Sim Frame** | *rend\_sim\_frame* – During simulation, after each new frame is simulated, this option will start a render with your currently assigned production renderer and with the render settings you have dialed in your Render Setup 3ds Max dialog. You need to have set a Render Output File in the 3ds Max Render Settings in advance. Phoenix will automatically append the frame number to the file name. You can use **Render On Each Sim Frame** when you leave a long simulation overnight, so you will have a ready rendered sequence when you come back to check on it, even if the entire simulation did not have time to finish. If you are doing test iterations and you don't need your fire/smoke to be rendered with production quality, you could enable the [GPU Preview extended viewport](#) and use the **Save Images During Sim** option instead.

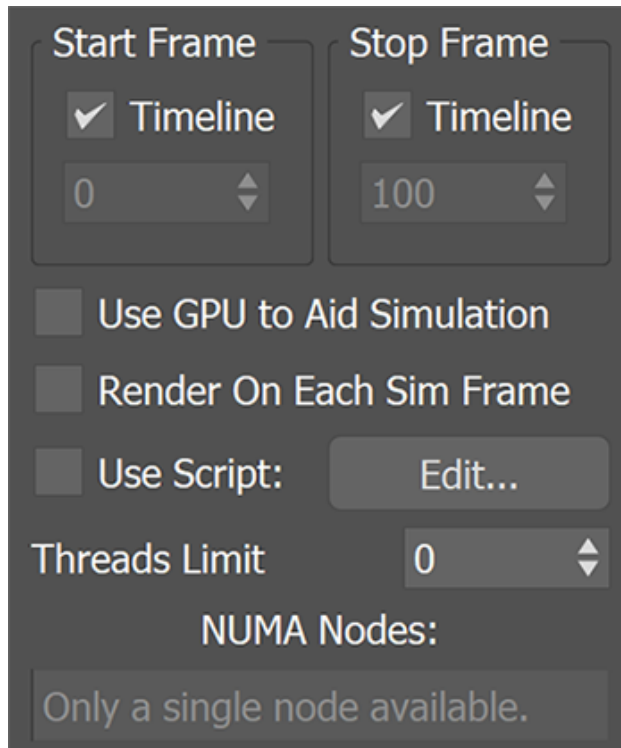
When doing test iterations on your simulation setup, you can also speed up the simulation by enabling **Cacheless Simulation** from the Output rollout. An alternative approach to **Render On Each Sim Frame** is to automatically start a sequence render after the entire simulation has finished by adding a ' *max quick render*' command under the *OnSimulationEnd()* callback of the **Use Script** option below.

**Use Script** | *script, use\_script* – Enables the use of a MAXScript during the simulation. Press **Edit...** to open a window with sample code and a short description of available commands. For more information, please see the [Phoenix MaxScript documentation](#).

**Threads Limit** | *maxthreads* – Specifies an upper limit on the number of threads used for the simulation. When the value is set to 0, the maximum number of threads (cores) will be used.

**NUMA Nodes** – Specifies which nodes will be used if the simulation is executed by a NUMA machine. Phoenix is very sensitive to the memory transfer and if all the nodes are used, performance might decrease due to low bandwidth between the different nodes.

NUMA stands for Non-Uniform Memory Access. It can be used to restrict the threads used for simulation based on the physical CPUs available on the system (in a multiprocessor system). In this way, better memory access can be achieved when multiple simulations are run on the same machine.



The screenshot shows a settings panel with a dark grey background. At the top, there are two sections: 'Start Frame' and 'Stop Frame'. Each section has a checked checkbox labeled 'Timeline' and a numeric input field. The 'Start Frame' input is set to '0' and the 'Stop Frame' input is set to '100'. Below these are three unchecked checkboxes: 'Use GPU to Aid Simulation', 'Render On Each Sim Frame', and 'Use Script:'. To the right of the 'Use Script:' checkbox is a button labeled 'Edit...'. Below these is a 'Threads Limit' label followed by a numeric input field set to '0'. At the bottom, there is a 'NUMA Nodes:' label followed by a text box containing the message 'Only a single node available.'

## Cache File Content

This area displays information on the current simulation. Any data from channels that have been loaded into the simulator are shown. This includes the minimum and maximum ranges for each channel. Check the Grid Channel Ranges and Particle Channel Ranges pages for more information.

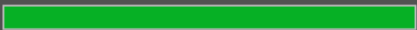
Note that the **Container Dimensions** show the loaded cache size as seen in the scene, in the currently selected units. This size is not multiplied by the [Scene Scale](#) parameter from the Grid rollout. If you want to see how big the container is as seen by the Phoenix simulator when the Scene Scale parameter is used, check the **X, Y, Z** info in the Grid rollout.

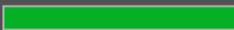
Total Elapsed 53.8s

Estimated Time Left

Frame Elapsed 1.3s

Performance 1.27M vox/s

Frame 

Total 

Container Dimensions (units):

105.103 x 102.767 x 242.904

Cache File Content:

[GRID: 1 647 360 cells, 90x88x208]  
- Liquid/Temperature (300.00 : 2000.00)  
- Velocity (0.03 : 739.37)  
\* Simulated with Phoenix 5.10.01