



NVIDIA[®] mental ray[®] Standalone

Frequently Asked Questions

Version 1.0

May 2014

Copyright © 2014 NVIDIA Corporation. All rights reserved.

This document is protected under copyright law. The contents of this document may not be translated, copied or duplicated in any form, in whole or in part, without the express written permission of NVIDIA Corporation.

The information contained in this document is subject to change without notice. NVIDIA Corporation and its employees shall not be responsible for incidental or consequential damages resulting from the use of this material or liable for technical or editorial omissions made herein.

NVIDIA, the NVIDIA logo, imatter, IndeX, Iray, mental images, mental ray, and RealityServer are trademarks and/or registered trademarks of NVIDIA Corporation. Other product names mentioned in this document may be trademarks or registered trademarks of their respective companies and are hereby acknowledged.

Table of Contents

1	How do I install Standalone ?.....	3
2	How do I activate Standalone ?.....	3
3	How do I use Standalone ?.....	3
4	How many licenses do I need for rendering with Standalone ?	3
5	How many licenses are needed to render on a quad-core machine ?	4
6	How do I render my Autodesk Maya scene with Standalone ?	4
7	How do I render my Autodesk 3ds Max scene with Standalone ?.....	5
8	I exported my scene as .mi file on Windows. How do I render it on Linux ?.....	5
9	I have an NVIDIA GPU that supports CUDA. Does Standalone make use of it ?.....	5

1 How do I install Standalone ?

After download of the installer package for your platform, start the installer executable with administrator privileges. It will guide you through the installation process of the software package. The Standalone executable with all its related plugins and shader packages, as well as documentation will be placed into the system specific application folder, unless you chose a custom location during the installation process. The Standalone program file is called "ray" and may be started immediately, like to check a working setup. However, it requires a license to be able to render a scene file, see next question how to *activate*.

2 How do I activate Standalone ?

Standalone needs a license to run. A separate license server has to be installed on a machine in your network that is visible to the computers used for rendering. The licenses itself can be purchased on www.nvidia-arc.com/mentalray. To be able to generate the actual licenses for that server, specific information about that machine will need to be provided. The issued licenses will be sent back by email. Once they are added to the license server database and the rendering machines are able to contact that server machine, the Standalone is ready to render.

3 How do I use Standalone ?

The Standalone is a command-line tool that reads scene files in the mental ray interchange proprietary .mi format (or receives streams in this format via input pipe). In order to render scenes from Applications like Autodesk Maya, 3ds Max, or Softimage, a .mi scene file representation of the application specific scene data needs to be created, typically by an "export" operation, like in Maya. Alternatively, embedded versions of mental ray may be used to "echo" the internal scene database to a .mi file, like in 3ds Max and Softimage. Such a .mi file contains references to external data, like texture files, custom shaders, or cache/assembly files. Some applications allow customizing how such linked content is written in the .mi file, with absolute, relative, or no path.

The Standalone can be configured to search in a list of user-defined directories for data of a certain type, so that the actual textures etc. do not need to be moved to a local subfolder in order to be found. However, keeping the project-related data in a common place can ease the managing of rendering jobs. A typical .mi file job for final rendering is self-contained, which means it describes all necessary render settings and comprises, or references, all required scene data. For larger jobs, the scene may be split into several .mi files that reference each other, either with a direct include or through on-demand loadable assembly files.

For the actual render job, almost all settings can be overridden on the command-line without the need to touch the content. That permits to optimize a Standalone rendering on a per job basis, for example to tune for certain hardware conditions, which is essential for render farm management software's job control.

4 How many licenses do I need for rendering with Standalone ?

The number of licenses will determine how many Standalone renderings can be executed simultaneously on a machine, or on several machines in your local environment ("render farm"). One such license is needed per running Standalone, no matter how the machine is equipped, like how

many CPUs or GPUs are available, what operating system is used, or if it is a 32bit or 64bit system. A running Standalone will keep its license checked out from the license server, taking away from the pool of available licenses.

For best performance on a single machine, only one Standalone process should run on it at any one time, because this one instance can leverage the full processing power and memory resources for rendering. On multi-core machines, mental ray will take advantage of the extra processors so that performance scales with the number of cores. Multiple Standalone processes can run simultaneously - each taking its own license - until the pool of licenses is fully used.

In the end, the number of licenses needed depends on the expected rendering throughput, which itself depends on the typical scene sizes to handle, the number of machines and their capacity and compute speed determining the average runtime, but also the exploited rendering effects and quality goals. For larger installations of rendering computer systems ("render farm") it is desirable to have as many licenses available as there are machines dedicated to rendering, to be able to fully load the system and take maximum benefit.

The number of licenses can grow with the requirements of the work, additional licenses can be purchased and added at any time.

5 How many licenses are needed to render on a quad-core machine ?

One license per render, eligible to use any number of CPUs or CPU Cores, or GPUs of the machine.

In other words, starting a single mental ray Standalone process will pull one license. A second one started in parallel on that same machine, or on a remote machine, will pull another license, independent of its hardware or operating system equipment.

6 How do I render my Autodesk Maya scene with Standalone ?

The most common way to render Maya scenes with Standalone is: by exporting the Maya scene, or parts of it, to a .mi file (*File > Export All*) and render it separately with Standalone on a dedicated machine.

This workflow allows greatest flexibility in terms of scheduling and resource planning of a separate rendering pipeline, like a render farm. The file export of a Maya scene can be done manually from within a running Maya, which would require a Maya license. It can also be done on the command-line using Maya's *Render* command, which does not require a Maya license. To be able to render any of the Maya specific shading and lighting effects, the Standalone needs to be configured to find and use Maya's implementation of these effects for mental ray, called "shaders". Adding the shader location to the search path for Standalone is sufficient, and doing that in the global configuration file "rayrc" will fix it for all uses of Standalone on that machine.

For this to work reliably and give consistent results, the version of Standalone should match the mental ray version embedded into Maya very closely. They do not need to be exactly identical, though, in normal cases.

Other external scene elements like textures and cache files may need to be grouped, and possibly copied to remote locations, so that Standalone will find them easily. Various Standalone command-line options help to adjust the search path for certain files, or file types, by adding custom file locations.

Maya's advanced scripting capabilities can be used to automate the .mi file export and start of the Standalone rendering in the background, and even make that a push of a button in the UI.

7 How do I render my Autodesk 3ds Max scene with Standalone ?

The most common way to render a 3ds Max scene with Standalone is: generate an .mi file by actually "rendering" the 3ds Max scene with mental ray set to "echo" its scene database to file, rather than rendering an image. To achieve this, select mental ray as the *Production Render*, in the *Render Setup* window, and go to *Processing / Translator Options / Export to .mi File* and set an output .mi file name, so that the option *Export on Render* gets enabled. Render normally.

To be able to render any of the 3ds Max specific shading and lighting effects, the Standalone needs to be configured to find and use the 3ds Max implementation of these effects for mental ray, called "shaders". Otherwise, it works similar to the Maya workflow (see above).

Please note, that 3ds Max is not available for Linux or Mac OS X, which typically prevents from rendering 3ds Max content on these platforms out-of-the-box. However, Autodesk does provide the required mental ray shaders for 3ds Max also built for these platforms with its own mental ray Standalone product. Those shader binaries can be used together with the NVIDIA mental ray Standalone without problems.

8 I exported my scene as .mi file on Windows. How do I render it on Linux ?

The export procedure in the applications will typically create references to shader or plugin binaries for the current platform, like "shader.dll" on Windows (likewise "shader.dylib" on Mac, and "shader.so" on Linux). The mental ray Standalone will handle such .mi files automatically, and correct any system dependent references to the correct file extension on the target rendering system. Other assets created on one platform but used on another, like texture files, should be saved to a platform independent format. Almost any common image file format is platform independent, with the exception of proprietary formats like mental ray's ".map" file variations. The mental ray specific files used to store Finalgather Map, or Photon Map, are proprietary but standardized, and are readable by mental ray on any other platform.

9 I have an NVIDIA GPU that supports CUDA. Does Standalone make use of it ?

Yes it will, if one of the GPU utilizing features is enabled for rendering. That currently includes: the Iray Rendering Mode, and computing Ambient Occlusion. If the detected GPU does not meet the requirements of the algorithm then the effect will be rendered on the CPU. More rendering effects are currently developed to take benefit of the GPU going forward.