



NVIDIA Iray Server

User and administrator guide

14 January 2026
2025.0.4



NVIDIA Iray Server – User and administrator guide

Copyright Information

© 2025 NVIDIA Corporation. All rights reserved.

Document build number rev387700.2953

Contents

Preface	1
Purpose of this document	1
Audience	1
How this document is organized	1
Part 1 Installing and starting Iray Server	3
1 Introducing Iray Server	5
1.1 Server-side rendering	5
1.1.1 Supported workflows	5
1.1.2 Supported rendering modes	5
1.1.3 Browser-based user interface	6
1.2 Supported configuration modes	6
1.2.1 The configuration modes	6
1.3 Licensing Iray Server	7
2 Installing Iray Server	9
2.1 System requirements	9
2.1.1 Supported operating systems	9
2.1.2 Hardware requirements	9
2.1.3 Software requirements	9
2.2 Configuration requirements	10
2.2.1 Setting up a master node	10
2.2.2 Performance considerations	10
2.3 Updates and upgrades	11
3 Starting Iray Server	13
3.1 Starting Iray Server under Windows and Linux	13
3.1.1 Starting Iray Server under Windows	13
3.1.2 Starting Iray Server under Linux	14
3.2 Logging in to the Iray Server GUI	14
3.2.1 Logging in	14
3.2.2 Recovering from log-in errors	15
3.2.3 Logging out	17
3.3 Managing your account	17
3.3.1 Checking your user profile	17
3.3.2 Changing your password	17
3.3.3 Enabling desktop notifications	18

3.4	The iray_server command and options	18
3.4.1	General options	18
3.4.2	Path options	19
3.4.3	HTTP options	19
3.4.4	HTTPS options	20
3.4.5	Cluster options	20
3.4.5.1	Establishing the master node of a cluster	21
3.4.5.2	Cluster formation with TCP/IP	21
3.4.5.3	Cluster formation with multicast	21
3.4.5.4	Iray Server on a single host	22
3.4.6	Farming options	22
3.4.7	Iray Bridge protocol options	23
3.4.7.1	Iray Worker configuration	23
3.4.7.2	Cache manager options	23
	Part 2 Running jobs	25
4	Touring the Queue page	27
5	Adding jobs to the queue	31
5.1	Specifying a gamma correction value	31
5.2	Submitting your job	31
5.3	Creating a job	32
5.4	Editing a queued job	32
5.5	A look at job parameters	33
5.6	Editing job parameters for stereo image pairs	34
5.7	Archiving jobs	35
6	Managing the queue	37
6.1	Stopping a job	37
6.2	Restarting a job	37
6.3	Moving a job in the queue	37
6.4	Deleting (removing) a job in the queue	38
7	Managing results	39
8	Managing the archive	41
	Part 3 Managing settings, resources, and users	43
9	Managing configuration settings	45
9.1	Touring the Settings page	45
9.2	Specifying general settings	45
9.3	Configuring results settings	46
9.3.1	Changing the global storage directory	46
9.3.2	Creating zip files for downloads	47
9.3.3	Specifying an update interval	47

9.3.4	Buffering interim results	47
9.4	Adding custom commands	47
10	Starting and stopping the render queue	49
10.1	Starting the render queue	49
10.2	Stopping the render queue	49
10.3	Switching between streaming and queuing modes	49
11	Managing rendering resources	51
11.1	Touring the Resources page	51
11.2	Checking the supported versions of Iray Bridge	51
11.3	Checking the cache manager settings	52
11.4	Checking the available storage for render results	52
11.5	Checking the server configuration	53
11.6	Checking and downloading the log of the master node	53
11.7	Benchmarks	54
12	Managing users	55
12.1	Touring the Users page	55
12.2	Adding a user	56
12.3	Deleting a user	56
12.4	Editing a user profile	57
13	Troubleshooting	59
13.1	Clustering	59
13.2	Farming	59
13.3	Out-of-memory	59



Preface

Purpose of this document

NVIDIA® Iray® Server rendering solution provides multi-user, high-performance, distributed rendering capabilities for Iray-compatible applications. This document describes Iray Server capabilities and how to install, configure, run, and manage it.

Audience

This document is intended for:

- Users of [Iray-compatible applications](#)¹
- Administrators who are responsible for installing Iray Server and managing settings, resources, and users

How this document is organized

This document is organized as follows:

- Part 1. Installing and starting Iray Server:
 - [Introducing Iray Server](#) (page 5) describes what Iray Server is and what it is used for.
 - [Installing Iray Server](#) (page 9) describes hardware, software, and configuration requirements as well as performance considerations.
 - [Starting Iray Server](#) (page 13) describes how to start Iray Server from the command line as well as from a Windows desktop icon. It also describes how to log in to the browser-based UI for Iray Server. From this UI you manage the render queue and render results.
- Part 2. Running jobs:
 - [Touring the Queue page](#) (page 27) describes the Queue page where you manage the job queue.
 - [Adding jobs to the queue](#) (page 31) describes the various ways you can add jobs to the queue.
 - [Managing the queue](#) (page 37) describes the various tasks for queue management.
 - [Managing results](#) (page 39) describes the Results page where you can view, download, and delete job results.
 - [Managing the archive](#) (page 41) describes how to view, delete, and add archived jobs to the render queue.

1. <https://www.irayplugins.com>

-
- Part 3. Managing settings, resources, and users:
 - [Managing configuration settings](#) (page 45) describes how to specify the configuration settings for your Iray Server setup.
 - [Starting and stopping the render queue](#) (page 49) describes how to start and stop the render queue and switch between streaming and queuing modes.
 - [Managing rendering resources](#) (page 51) describes how to manage rendering resources.
 - [Managing users](#) (page 55) describes how to manage user access to Iray Server.
 - [Troubleshooting](#) (page 59) describes how to resolve errors that may arise when running Iray Server.

Note: Administrator rights are required to perform the tasks in Part 3.

Part 1 Installing and starting Iray Server

1 Introducing Iray Server

This chapter describes new rendering workflows supported by the NVIDIA® Iray® Server rendering solution, hardware and software requirements, and a link for licensing information.

1.1 Server-side rendering

The following sections describe the rendering workflows, supported Iray rendering modes, and the browser-based GUI for managing rendering requests when using Iray Server.

1.1.1 Supported workflows

Iray Server is a multi-user, high-performance, distributed rendering solution for [Iray-compatible applications](#).¹ It supports server-side rendering for:

Queuing

Iray Server renders offline using batch processing. Client-side users access render results from Iray Server for viewing and downloading.

Streaming

Iray Server renders on an external host while client-side users interact with the streamed render results. Interactive changes to data on the client-side are uploaded to the server for rendering updates. Only images are sent to the client in this workflow.

Any host running Iray Server can process jobs sent to it for offline rendering. The host that clients connect to and send jobs to is termed the *master node* while all other hosts in the cluster are termed *rendering nodes*. In practice, the client (plug-in) first connects to the master node's address to verify the user has access to Iray Server. Sending a job "to the queue" outputs the required Iray render data to the Iray Server master node as a job for processing. Sending a job is usually quite quick, after which the client application is free to do other work. Any number of jobs can be sent to Iray Server before they are processed, without any performance impact to the master node. Submitted jobs can be managed and modified from the Iray Server GUI.

1.1.2 Supported rendering modes

Iray Server supports the following NVIDIA® Iray® rendering modes:

Iray Photoreal

Generates physically-based photorealistic imagery. Images are progressively refined to provide full global illumination including caustics, sun studies, and luminance distributions.

1. <https://www.irayplugins.com>

Iray Interactive

Targets a look which is consistent with the physically-based results of Iray Photoreal but uses faster, less accurate rendering algorithms than Iray Photoreal. The default settings target a look that is very close to Iray Photoreal, but global illumination effects and ambient occlusion are not enabled.

See the [Iray Plugins page](#)² for a list of Iray-compatible applications that support these render modes.

1.1.3 Browser-based user interface

To manage rendering requests, Iray Server offers a browser-based interface, which is supported by a built-in HTTP server:

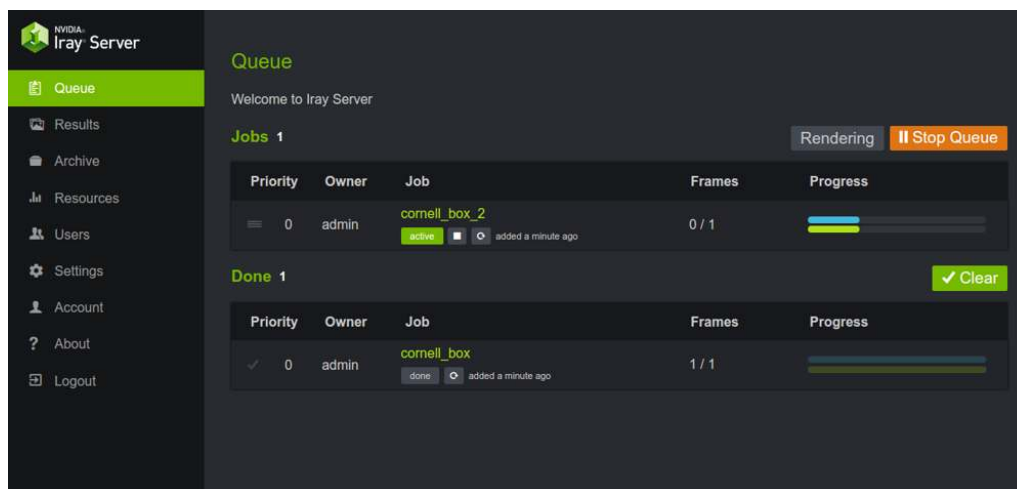


Fig. 1.1 - The Iray Server user interface

The user interface is minimalistic in design. The left panel displays the GUI pages with active links. The right panel displays the content of the current page. Secondary functionality is accessed from links and switches. Users can inspect the render queue, change the order of the jobs in the queue, view, archive, and download render results. Users with administrator rights can manage configuration settings, resources, and users. See Parts 2 and 3 of this document for a detailed description of the Iray Server GUI and supported tasks.

1.2 Supported configuration modes

The following sections describe the configuration modes supported by Iray Server.

1.2.1 The configuration modes

Iray Server supports the following configuration modes:

Standalone mode

Iray Server runs on a [single host](#) (page 22). You can install Iray Server on:

- The same host as your Iray-compatible application
- A separate host

2. <https://www.irayplugins.com>

Cluster mode

Multiple Iray Server instances are [connected to each other](#) (page 20) in a local area network (LAN). Hosts in the LAN can be enabled or disabled to participate in rendering. Iray Server supports UDP multicast and TCP/IP with discovery.

Farming mode

Intended for larger rendering farms, where regular network rendering does not scale efficiently. Farming mode is typically used in digital special effects and animation production.

Farming mode uses the [same host configuration](#) (page 22) as cluster mode. The key difference lies in the rendering strategy for each frame:

- In *cluster mode*, the queue is processed frame by frame. All render nodes in the cluster work on the same frame.
- In *farming mode*, each render node receives a single frame. This reduces the load on the network and enables multiple frames to be processed simultaneously.

Assuming a minimum of 24 frames per second, incremental speed ups per frame can have a significant cumulative impact over time. Time saved in rendering can be used to improve the quality of scene elements or lighting where that makes sense. For example, using farming mode for the quick evaluation of scenes during the early stages of scene development, enables deficiencies to be addressed when it is both the most efficient and the least expensive time to do so.

You specify a configuration mode when you start Iray Server.

1.3 Licensing Iray Server

See the [Iray Server³](#) web page for details about trying and licensing Iray Server.

3. <https://www.irayplugins.com/iray-server/>

2 Installing Iray Server

This chapter describes the requirements for installing Iray Server under Windows and Linux operating systems and important performance considerations for cluster rendering.

2.1 System requirements

The following sections describe system requirements for Iray Server.

2.1.1 Supported operating systems

Iray Server can be installed under Windows or Linux. Note the following:

For Windows 10, 64-bit:

An installer is provided. By default, Iray Server is installed in:

`C:\Program Files\NVIDIA Corporation\Iray Server`

After Iray Server is installed, you can copy the Iray Server folder from Program Files to a different location. When you start `iray_server.exe` directly from the new location, subfolders are created in the new location for caching and image results.

The default cache and image results directory is:

`C:\ProgramData\NVIDIA Corporation\Iray Server`

You can start Iray Server from the icon installed on your desktop.

For Linux, 64-bit:

A gzipped tar file is provided, which you can unpack using the `tar xvfz` command.

2.1.2 Hardware requirements

Iray Server has the following GPU and CPU requirements:

GPU acceleration:

Maxwell or later NVIDIA GPUs

Streaming:

Professional GPUs — Quadro, Tesla, or Grid — are required on the server side

CPUs

x86-64 processors

2.1.3 Software requirements

Iray Server is a companion application for [Iray-compatible applications](#).¹ Note the following:

1. <https://www.irayplugins.com>

- No application other than Iray Server is required on dedicated rendering hosts.
- Ensure that the Iray Bridge versions that this release of Iray Server supports are supported by your Iray-compatible applications.

2.2 Configuration requirements

Iray Server is designed to support a range of [configurations](#) (page 6) with minimal restrictions:

- Heterogeneous client-server configurations are supported, for example, a Windows client and Linux rendering nodes.
- Heterogeneous server configurations are not supported. Every host in an Iray Server cluster must run under the same operating system.

At least 100-Mbit/s Ethernet is recommended between clustered hosts. Between the client and a streaming server, 10-Mbit/s is acceptable.

2.2.1 Setting up a master node

When configuring Iray Server for distributed rendering, only the *master node* can be used to receive jobs from client applications. The master node also provides the Web UI for Iray Server users. You specify the [master node](#) (page 21) from the command line in a terminal window.

Note: All other nodes in an Iray Server cluster are rendering nodes only.

When a job is added to the queue, the required Iray rendering data is sent to the Iray Server master node as a set of instructions for processing. Any number of jobs can be sent to Iray Server without a performance impact on the master node. Submitted jobs are managed and modified from the [Queue](#) (page 27) page.

2.2.2 Performance considerations

Rendering is most efficient when the performance of all hosts in a cluster is similar. However, you can mix hosts with different performance levels with fairly minimal impact. Slow hosts will not hold back much faster hosts in the same cluster; faster hosts will simply contribute more iterations per minute than the slower hosts.

Hosts added to the cluster after queue processing has begun will be included for use on the next frame. You can remove hosts from the cluster without causing errors in an active job, other than losing the iteration a particular host is currently rendering. A lost iteration is added to the next set of iterations sent to the master node.

An Iray Server cluster runs independently of any other Iray Server. If you have a large number of potential render farm hosts, the decision about how large a cluster should be is usually one of efficiency and rendering job type. Tests show that more than 70 percent efficiency is possible with 20 similar hosts running a complex render job. Very fast render jobs (those completing under a minute) may see less than 70 percent efficiency with just a few hosts in the cluster.

2.3 Updates and upgrades

When you update or upgrade Iray Server, passwords, the cache, and render results remain unchanged.

3 Starting Iray Server

This chapter describes how to start Iray Server under Windows and Linux, how to log in to the Iray Server GUI, and how to manage your user account settings. The chapter concludes with an extensive description of the `iray_server` command-line options.

3.1 Starting Iray Server under Windows and Linux

The following sections describe how to start Iray Server under Windows and Linux.

3.1.1 Starting Iray Server under Windows

To start Iray Server under Windows:

1. Double-click the Iray Server icon on your desktop. A terminal window is opened and Iray Server is started.

```
[Wed, 06 Jan 2016 15:02:56] Configuring database...
[Wed, 06 Jan 2016 15:02:56] Database schema generated
[Wed, 06 Jan 2016 15:02:56] Iray Server starting...
[Wed, 06 Jan 2016 15:02:56]   HTTP Port:          9090
[Wed, 06 Jan 2016 15:02:56]   Install Path:       C:\Program Files\NVIDIA Corporation\Iray Server 2015
[Wed, 06 Jan 2016 15:02:56]   Disk Cache Location: path:cache
[Wed, 06 Jan 2016 15:02:56]   Log File:           iray_server.log
[Wed, 06 Jan 2016 15:02:56]   Plugin Search Path:
[Wed, 06 Jan 2016 15:02:56]   Multicast Address:
[Wed, 06 Jan 2016 15:02:56]   Multicast Increment: 0
[Wed, 06 Jan 2016 15:02:56]   TCP/IP Discovery Address:
[Wed, 06 Jan 2016 15:02:56]   Cluster Interface:
[Wed, 06 Jan 2016 15:02:56]   Head Node Address:
[Wed, 06 Jan 2016 15:02:56] Loading neuray...
[Wed, 06 Jan 2016 15:02:56] Loaded C:\Program Files\NVIDIA Corporation\Iray Server 2015\libneuray.dll
[Wed, 06 Jan 2016 15:02:56] Using license file: C
[Wed, 06 Jan 2016 15:02:56] [Streaming] session state callback created
[Wed, 06 Jan 2016 15:02:56] Configuring network...
[Wed, 06 Jan 2016 15:02:56] Cluster manager connection enabled
[Wed, 06 Jan 2016 15:02:56] Connecting to cluster manager 'wss://8.36.120.201:443' with user 'Robert'
[Wed, 06 Jan 2016 15:02:58] Connected to cluster manager
[Wed, 06 Jan 2016 15:02:58] [Mode] Switching mode to Standalone
[Wed, 06 Jan 2016 15:02:58] Found pro GPU: [2] Quadro K6000
[Wed, 06 Jan 2016 15:02:58] Pro GPU is available (Quadro or Tesla)
[Wed, 06 Jan 2016 15:02:58] Streaming is supported: yes
[Wed, 06 Jan 2016 15:02:58] Local rendering is enabled
[Wed, 06 Jan 2016 15:02:58] Setting disk cache location to path:cache
[Wed, 06 Jan 2016 15:02:58] Loading plugins...
[Wed, 06 Jan 2016 15:02:58] Plugins loaded
[Wed, 06 Jan 2016 15:02:58] Starting neuray...
[Wed, 06 Jan 2016 15:02:58] [Cluster_callback] this host has the id 1
[Wed, 06 Jan 2016 15:02:58] Starting bridge...
[Wed, 06 Jan 2016 15:02:58] Enable streaming...
[Wed, 06 Jan 2016 15:02:58] Created bridge application for streaming at /bridge
[Wed, 06 Jan 2016 15:02:58] Streaming has been enabled
[Wed, 06 Jan 2016 15:02:58] Created bridge application at /iray_server_bridge.ws
[Wed, 06 Jan 2016 15:02:58] Bridge protocol version: 246000.2384
[Wed, 06 Jan 2016 15:02:58] -----
[Wed, 06 Jan 2016 15:02:58] Using trial license: 26 days and 8 hours left
[Wed, 06 Jan 2016 15:02:58] -----
[Wed, 06 Jan 2016 15:02:58] Starting http server...
[Wed, 06 Jan 2016 15:02:58] Available network interfaces:
[Wed, 06 Jan 2016 15:02:58]   10.21.120.154   Intel(R) Ethernet Connection (2) I218-U
[Wed, 06 Jan 2016 15:02:58] Iray Server started - listening on all network interfaces on port 9090
```

Fig. 3.1 - Iray Server running in a terminal window

If Iray Server fails to start, go to the terminal window and check the log for errors.

2. From a browser, enter the appropriate URL for your configuration. The login page for Iray Server is displayed.

You can configure command-line options by editing the settings of the Iray Server desktop icon. Alternatively, you can configure [command-line options](#) (page 18) when you start Iray Server directly from the Windows command line.

3.1.2 Starting Iray Server under Linux

To start Iray Server under Linux:

1. Open a terminal window.
2. From the command line, enter `iray_server` and the appropriate command-line options for your [configuration mode](#) (page 6).

The following console message is displayed:

```
Iray Server is starting ...
```

3. To verify your configuration mode, check for the Mode message, for example:

```
[Mode] Switching mode to Standalone
```

4. To verify that the start-up process is successfully completed, check for the following console message:

```
Iray Server started - listening on [port]
```

The Resources page also displays the current configuration mode as well as the type of configuration.

3.2 Logging in to the Iray Server GUI

The following sections describe how to log in and log out of Iray Server GUI.

3.2.1 Logging in

To log in to the Iray Server GUI:

1. From a web browser, enter the appropriate URL for your Iray Server configuration. The Iray Server Sign In page is displayed:




Fig. 3.2 - The Iray Server Sign In window

After installation, the user name and password are admin. When you log in for the first time, you are asked to change your password.

2. Type your name and password and click Login. The [Queue page](#) (page 27) is displayed:

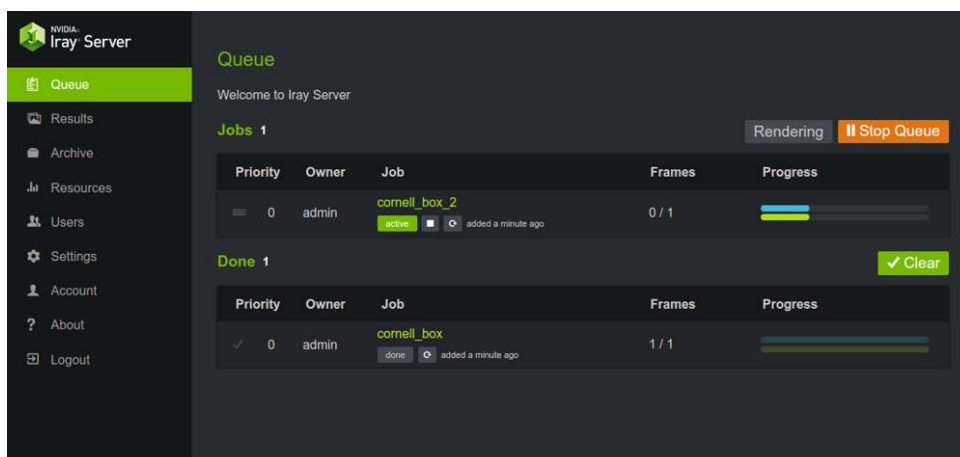


Fig. 3.3 - The Queue page

You manage your jobs from this page.

3.2.2 Recovering from log-in errors

When you attempt to log in to a cluster, you may accidentally attempt to log in to a render node rather than a master node. Depending on the error, one of the following warning messages are displayed together with information about what to do next:

A warning that contains the address of the master node

If you log in to a render node, this warning is displayed:




Fig. 3.4 - Continue to master

Click Continue. You are re-directed to the master node where you can continue the log-in operation.

A warning that the cluster has no master

If a cluster has no master node, this warning is displayed:




Fig. 3.5 - Waiting for master

Start a master server before continuing the log-in operation.

A warning that there are multiple masters

In the rare case that you try to log in to a cluster with two masters, the following warning is displayed:




Fig. 3.6 - Multiple masters

Restart the cluster with one master only before continuing the log-in operation.

3.2.3 Logging out

To log out:

1. Click Logout in the navigation panel. A confirmation window is displayed:




Fig. 3.7 - The log-out confirmation window

2. Click Logout. If the queue is running, any waiting jobs will continue to be rendered.

3.3 Managing your account

The following sections describe how to manage your user account. To manage the accounts of other users, you need administrator rights.

3.3.1 Checking your user profile

To check your user profile, open the Account page.




Fig. 3.8 - The Account page

You can:

- Change your [password](#) (page 17)
- Enable [desktop notifications](#) (page 18)

You cannot edit any other fields displayed in your profile.

3.3.2 Changing your password

To change your password:

1. Click Account from the navigation panel. The Account page is displayed.
2. Under Settings, click Change Password. The Change Password window is displayed.

3. Type your current password, your new password, confirm, and click **Save**.

3.3.3 Enabling desktop notifications

Note: Desktop notifications are not supported by all browsers.

To enable desktop notifications:

1. Click **Account** from the navigation panel. The **Account** page is displayed:

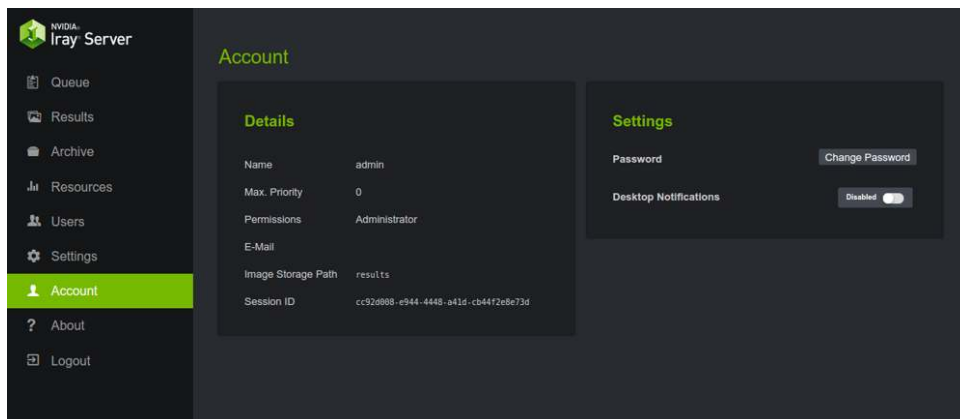


Fig. 3.9 - The Account page

2. Under **Settings**, check the status of **Desktop notifications**. If disabled, click the button to enable it. If your browser does not support desktop notifications, an information message is displayed and this feature remains disabled.

Desktop notifications are sent to you when a job is completed. The first time that a notification is sent to you from Iray Server, a pop-up window is displayed. Click **Allow** if you want to receive desktop notifications.

3.4 The iray_server command and options

You start Iray Server from the command line:

```
iray_server
```

Starts Iray Server

```
iray_server -help
```

Displays a list of command options

The following sections provide a comprehensive list of `iray_server` options. Most options described in the following sections are supported by all [configuration modes](#) (page 6).

3.4.1 General options

```
-v [-version
```

Displays the version and exits.

```
-h [-help
```

Displays the help and exits.

-rlm-activate

Specifies the activation key to check out an RLM retail license. After the key is successfully activated, you do not need to re-enter it.

Unix users: On start-up, the following message is displayed:

```
[RLM] Warning: Could not create directory to store rlm license files
in: "/var/opt/migenius/Iray Server/license/"
```

Root permission is required to write to this location. Alternatively, perform the following steps to enable Iray Server to find the license server:

1. Create a license file that points to the license server. This license file should contain one line: `host hostname`.
2. Point the variable `migenius_LICENSE` or the variable `RLM_LICENSE` to the license file.

-rlm-revoke

Revokes the active RLM retail license.

-start-queue

Starts Iray Server with the queue running. If this option is not supplied, the queue will accept jobs but will not render them until you start the queue using the browser-based UI.

-open-browser

At startup, the web interface of Iray Server is opened in the browser.

3.4.2 Path options

-install-path

Installation path containing the `html` folder and the `benchmarks` folder. The default is the local working directory.

-flexnet-path

Sets the expected FlexNet license path or server address.

-plugin-path

Sets the plugin search path.

-log-file

Write log messages to this file. If empty, file logging will be disabled.

3.4.3 HTTP options

-p [-port

HTTP port.

-l [-listen

Bind HTTP server to a specific IP address. Changing the setting is only required in rare cases for hosts with multiple network interfaces.

-connection_limit

The maximum number of concurrent connections accepted by the server. There is no predefined upper limit.

3.4.4 HTTPS options

To set up Iray Server to use the HTTPS protocol, specify the following command options:

-https-port

HTTPS port. By default, the HTTPS port is not used.

-ssl-certificate

For the part of the argument *.cert*, you can specify the full path name.

An SSL certificate is required when you use the HTTPS protocol. Use the `openssl` command to generate a certificate as described below.

-ssl-key

For *iray_server.key*, you can specify the full path name.

An SSL private key is required when you use the HTTPS protocol. Use the `openssl` command to generate a private key as described below.

Generating private keys and certificates:

- To generate a private SSL key with a passphrase, enter:
`openssl genrsa -out iray_server.key 2048`
- To generate a matching SSL certificate, enter:
`openssl req -new -x509 -key iray_server.key -out iray_server.crt -days 365`

3.4.5 Cluster options

To build a cluster and coordinate the workload of the hosts, Iray Server can use either UDP with multicast or TCP/IP.

-c [-cluster]

The mode of the cluster, either OFF, TCP, or UDP. OFF means that Iray Server runs in standalone mode. TCP use TCP/IP unicast to send data between hosts. Specify `--interface-address` and `--master-address` (not needed for the master node) to use this mode. UDP use UDP multicast and unicast to send data between hosts. Specify `--interface-address` and `--multicast-address` to use this mode.

-master-address

Master address and port. In TCP mode, all nodes — except the master — must specify this address; the address is the same as the `--interface-address` option configured for the master node.

-multicast-address

Multicast address and port. In UPD mode, all nodes must specify the same multicast address and port and will then connect to each other automatically. The address should be

from the organization's local scope (239.0.0.0-239.255.255.255). The default 239.0.0.1:11000 will work fine in most cases. If more than one cluster is needed, then choose a different address for each cluster. Note that the chosen port is used as a base; each worker will use an increment of this port.

-interface-address

Sets the interface address to be used for unicast data. Used in both TCP and UDP modes. This address must be set to a specific address and port; "any" address (0.0.0.0) will not work. Note that the chosen port is used as a base; each worker will use an increment of this port.

-m [-master

Starts Iray Server as master / with web interface to manage jobs.

Other Iray Server nodes that were not started as master will participate in rendering but the database is stored on the master server. You can have only one master server in a cluster.

3.4.5.1 Establishing the master node of a cluster

The --master flag makes that host the master node. The master node hosts the Queue Manager interface, maintains the cache of submitted jobs, and is the default location for image results.

Each cluster must have a master to operate, and each cluster can only have one master. If a cluster is formed without a master, the resulting pool of hosts will wait for a master to join that cluster. If a host is started with --master and there is a master already, Iray Server cannot be used until only one master remains.

3.4.5.2 Cluster formation with TCP/IP

Iray Server will use the TCP/IP protocol to form a cluster when the IP address of the master node is provided as a startup argument along with the --cluster argument (without a number).

For example, assuming that your master node has an IP address of 192.168.2.1, your startup arguments for forming a cluster would be:

<code>iray_server --cluster 192.168.2.1 --master</code>	<i>On the master node</i>
<code>iray_server --cluster 192.168.2.1</code>	<i>On all rendering nodes</i>

Note: There is no cluster number with TCP/IP since the cluster is explicitly defined by the master's address.

3.4.5.3 Cluster formation with multicast

Iray Server uses the multicast protocol to automatically locate hosts with the same cluster number. Nothing more than the -c# options needs to be supplied for this to work. The -c option, followed by an integer argument, tells the host what cluster number to join.

For example, to create cluster 1 :

```
iray_server -c1 --master      On the master node
iray_server -c1              On all rendering nodes
```

While multicast is quite easy to setup, some networks may constrain multicast bandwidth, resulting in greater than normal file transfer times and poorer rendering efficiency. There also may be situations where different parts of a facility are on different multicast networks and so cannot see one another or where multicast is disabled in switches and routers. If you find one of these to be the case on your network, then you should probably switch to TCP/IP.

3.4.5.4 Iray Server on a single host

When Iray Server is started without any cluster arguments, it defaults to running individually. As a cluster of 1 (one), it is automatically its own master node, so there's no need to supply a `--master` argument.

In this mode, Iray Server can be used for either streaming or queuing. If streaming, it will still accept jobs for later processing in its queue. This mode will also work on the same host as an Iray Client to provide a background rendering capability. For example, you could be working with an Iray plugin, and output any number of jobs, keep working and then process them as you take a break from Iray authoring. If your host has multiple GPUs, you may want to assign one for interactive authoring and others to process the Queue as you work.

3.4.6 Farming options

In cluster mode, all render nodes in the cluster work on the same frame. If you want render nodes to work on different frames, as is the case in a standard render farm, you need to specify the `--farming` option in your `iray_server` command.

In addition, you must also specify the following command options:

- `--cluster` option to build a cluster and coordinate the machines' workloads. Iray Server can use either UDP with multicast or TCP/IP.

For more information about the cluster options, see the section [Cluster options](#) (page 20).

- `--master` option starts Iray Server as the master with a web interface to manage jobs.

For more information about the `--master` flag, see the section [Establishing the master node of a cluster](#) (page 21).

For example, assuming that your master node has an IP address of 192.168.2.1, your startup arguments for forming a server farm with TCP/IP would be:

```
iray_server --master --farming --cluster 192.168.2.1      On the master node
iray_server --farming --cluster 192.168.2.1              On all render nodes
```

If you experience unexpected behavior when rendering, check the section on [troubleshooting](#) (page 59). It may provide hints about the source of the problem and how to fix it.

3.4.7 Iray Bridge protocol options

Iray Server Version 1.2 and later support multiple versions of the Iray Bridge protocol. This means that Iray Server can work with client applications that use different versions of Iray.

Note:

- To verify the supported versions of Iray Bridge for your system, see [Checking the supported versions of Iray Bridge](#) (page 51).
- When using versions of Iray Server prior to 1.2, Iray Server and Iray-compatible applications must use the same version of the Iray Bridge in order to work together.
- Iray Server supports Iray version 2016.1 and later.

3.4.7.1 Iray Worker configuration

Configuration: To specify which Iray Workers to start, use the following command-line option:

`-workers`

Example: `--workers 334300.2780`. If not set, all workers for the particular version of Iray are started.

Ports: The ports for each Iray Worker are derived from the ports given to Iray Server by command-line options. Iray Server increments these port numbers for each Iray Worker so they do not collide with each other.

Iray Worker recovery: When an Iray Worker fails:

1. The current job is set to failed and must be restarted manually.
2. Iray Server automatically tries to restart the Iray Worker before processing subsequent jobs (if any) in the queue.

Note:

- In farming mode, Iray Worker recovery is supported for the master node only.
- Iray Worker recovery is supported for Iray Server versions 3.3 and later.

3.4.7.2 Cache manager options

`-cm-address`

The address to listen on. Do not change it unless there is a conflict in the port usage.

`-cm-path`

The directory to be used by the cache manager to store its data.

`-cm-low-water`

When memory usage (in MB) drops below this value, garbage collection is ended.

-cm-high-water

When memory usage (in MB) exceeds this value, garbage collection is started. To disable garbage collection, set the value to 0.

Part 2 Running jobs

4 Touring the Queue page

This chapter describes the Queue page where you can create, submit, and archive jobs, and manage the job queue:




Fig. 4.1 - The Queue page. By default, your view of the queue is limited to your jobs. To view all jobs in the queue and perform other administrator tasks, you need administrator rights.

The following tables are displayed on the Queue page:

Jobs

This table lists the jobs, if any, in the queue. However The table does not indicate which Iray Worker (version of Iray) is assigned to run a particular job.

The Jobs table includes a mini-toolbar:

Render | Awaiting jobs | Stopped

The current status of the queue is displayed.

Start Queue | Stop Queue

Click this button to start or stop the queue. This action requires administrator rights.

Done

This table lists finished jobs. This list is intended to provide an overview of recently completed jobs. For a complete list of job results, see the table on the [Results](#) (page 39) page.

The Done table includes the Clear action button. Click this button to [archive](#) (page 35) the jobs and clear the table.

The following information is displayed for each job listed in Jobs and Done tables:

Priority

The [priority level](#) (page 37), which determines the order in which jobs are run.

Owner

The user ID of the person who submitted the job.

Job

The name of the job. This is the name assigned by the person who submitted the job.

When you click on the name, a description of the job is displayed:

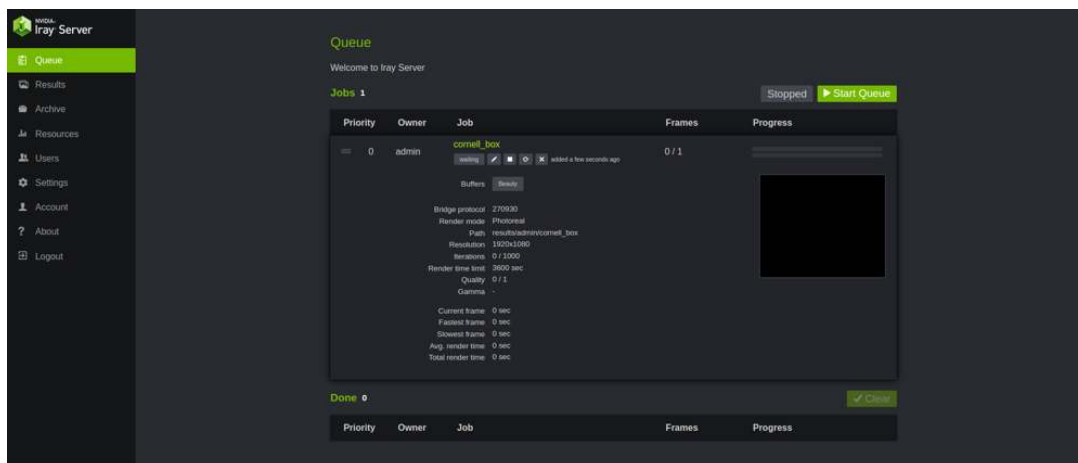


Fig. 4.2 - A job description is displayed when you click on the job name

Frames

The number of frames generated compared to the number of frames submitted.

Progress

A blue progress bar is displayed for each host used to render a particular job. The bar shows the rendering status for a particular frame. For animations, the green bar represents the percentage of the generated animation that is completed. Progress calculations are based on user-specified, [editable](#) (page 32) completion criteria for iterations, render time (hh:mm:ss), and quality, and which criterion is expected to be satisfied first.

Figure 4.3 shows the status of five hosts configured for farming mode, where multiple frames are rendered simultaneously:





Fig. 4.3 - Tracking the progress for a job rendered on multiple hosts


Figure 4.4 (page 29) is captured from the same five host configuration in farming mode. You can see that the first job is almost complete, leaving two hosts idle. The master server automatically assigns the idle hosts to the next job in the queue:



Priority	Owner	Job	Frames	Progress
2	admin	menger3-2 active added 4 minutes ago	33 / 36	Progress bar (33/36)
2	admin	menger3-3 active added 4 minutes ago	0 / 36	Progress bar (0/36) tray0 18% tray1 76% 2%

Fig. 4.4 - Tracking the progress of two jobs rendering on multiple hosts

Figure 4.5 shows the toolbar displayed with each job:



Priority	Owner	Job	Frames	Progress
50	robert	Diamonds (5) stopped added 6 minutes ago	10 / 60	Progress bar (10/60)

Fig. 4.5 - The toolbar

The toolbar provides the following status information and action buttons:

Stopped | Waiting | Running

The current status of the job is displayed.

Edit

Click this button to display the [Edit](#) (page 32) page where you can change the job parameters.

Start | Stop

Click this toggle button to [start](#) (page 37) or [stop](#) (page 37) the job.

Copy

Click this button to [copy](#) (page 32) the job, edit its parameters, and add it to the queue.

Remove

Click this button to [remove](#) (page 38) the job from the queue. You can add the job to the queue later on.

Note that removing and deleting jobs are not the same thing. When you delete a job, you can no longer retrieve it.

Time

The amount of time that has elapsed since the job was added to the queue is displayed. When two jobs have the same priority level, elapsed time is used to determine priority. In this case, the job that has been waiting longest is run first.

5 Adding jobs to the queue

This chapter describes how you can add a job to the Iray Server queue.

You can add a job to the Iray Server queue as follows:

- By submitting a job from your [Iray-compatible application](#).¹
- By copying a submitted job, editing it, and adding it to the queue. You can make a copy of a job that is waiting to run, that is finished running, or archived.

This section also discusses gamma correction and makes recommendations concerning gamma correction values for your jobs.

5.1 Specifying a gamma correction value

Gamma correction is represented by a floating point value that describes a non-linear function intended to correct for the non-linear response of display devices. The value you choose is dependent on the dynamic range of your image data:

- For low dynamic range file formats such as PNG and JPG, specify 2.2 as the gamma correction value.

The pixel values will be saved with the gamma correction applied (baked into the image). Image viewing applications expect a baked-in gamma correction value of 2.2 and will heighten the brightness level of darker pixels accordingly.

If you use the default gamma correction value (1.0), no gamma correction value is baked in by Iray Server. In image viewing applications, you can expect artifacts and detail in darker regions to be obscured.

- For high dynamic range file formats such as EXR, specify 1.0 (the default) as the gamma correction value.

Iray Server will not modify the image data when the default value (1.0) is used. Image viewing applications apply the gamma correction. This handling of gamma correction is particularly suited for post processing tasks where high dynamic range images are used for compositing or texture creation or modification.

5.2 Submitting your job

To submit your job:

1. Ensure that the [gamma correction value](#) (page 31) is appropriate.

1. <https://www.irayplugins.com>

2. Submit your job. Refer to the documentation for your [Iray-compatible application](#)² for instructions.
3. From the Iray Server [Queue page](#) (page 27), check the Jobs list to ensure that your job was submitted successfully.

5.3 Creating a job

From Iray Server, you can create a new job by making a copy of an existing job, editing its parameters to suit your needs, and adding it to the queue.

To create a new job:

1. Create a copy of an existing job.

From the Queue or Archive pages, select the job you want to copy and click **Copy**. You can select any job waiting in the queue to run or any job that has completed running.

The Copy page is displayed with a list of editable [job parameters](#) (page 33).

2. Modify the job parameters to suit your needs and click **Add to Queue**.

5.4 Editing a queued job

To edit a job in the queue:

1. From the Queue page, find the job that you want to edit and click the **Edit** icon. A job must be in **stopped** or **waiting** mode before you can edit it.

The Edit window is displayed with a list of [editable parameters](#) (page 33).

2. Modify the parameters to suit your needs and click **Save**.

A job may continue to move up the queue during the edit operation but it cannot be run. After you close the Edit window, the previous status of the job is restored — **waiting** or **stopped**.

2. <https://www.irayplugins.com>

5.5 A look at job parameters

The Copy and Edit pages display the following editable parameters when creating a job or modifying an existing job:

The screenshot shows a dark-themed interface for job parameters. At the top left, a green pencil icon and the word 'Copy' in green indicate edit mode. The parameters are as follows:

- Name:** TestProject (with a green checkmark and 'Available' text)
- Priority:** 5 (with a note: 'Lower values mean higher priority')
- Render mode:** Stereo Photoreal (dropdown menu)
- Stereo Pairing:** Left over Right (dropdown menu)
- Eye Separation Distance:** 0.314 (with a note: 'In scene units')
- Resolution:** 800 x 600 (with a lock icon)
- Buffers:** A table with columns: Buffer, Format, Bit Depth, Remove. One row shows 'Beauty', 'PNG', 'RGBA', and a 'Remove' button.
- Light Path Expressions:** A section with a dropdown arrow and the text 'Add light path expressions'. It contains three rows:
 - Max. iterations:** 100 (with a note: '1 - infinity')
 - Max. render time:** 60 (with a note: 'In seconds')
 - Max. quality:** 1 (with a note: '0 - infinity')

At the bottom right, there are 'Cancel' and 'Save' buttons.

Fig. 5.1 - The job parameter list. The page title identifies whether you are in Copy or Edit mode

Following is a description of the job parameters:

Name

The name of the new job, which must be unique. If the frame of the entry box is red, the displayed name is already used and must be changed before you can submit the job. If the frame of the entry field is green, the displayed name is unique.

Priority

A job's position in the queue is determined by its priority level.

Zero (0) is the highest priority level. You cannot assign a priority level to a job that is higher than the priority level assigned to you by the administrator.

Render mode

Supported render modes are Photoreal, Interactive, Stereo Photoreal, and Stereo Interactive.



Fig. 5.2 - Render modes dropdown menu

If you select a stereo render mode, [parameters](#) (page 34) controlling stereo image creation are displayed.

Resolution

Enter the resolution. You can unlock the ratio between height and width by clicking on the little lock symbol.

Buffers

A list of all buffers to be rendered. For each buffer the buffer type, the file format and the bit depth is specified. To remove buffers in the list, click Remove.

You can:

- Change the file format and the bit depth for each buffer in the list.
- Add new buffers by selecting a new buffer type from the drop-down menu to the left of the Add Buffer button. You must select an image file format and a bit depth for the output data. You can only select supported combinations of buffer type, image file format, and bit depth.

Light Path Expressions

Light Path Expressions (LPEs) enable you to render any interaction of light with any object into separate images, which can then be composited. To generate output files using LPEs, you specify a name, light path expression (LPE), format, and bit depth, and click Add. For more information about creating and using LPEs, refer to the documentation for your Iray-compatible application and the iray dev blog entry, [Compositing with Light Path Expressions](#).³

Max. iterations / Max. render time / Max. quality

Calculations for determining the end of the render operation can include any combination of these three criteria.

5.6 Editing job parameters for stereo image pairs

The following editable parameters are provided for setting up stereo image pairs:

Stereo pairing

Separate

Two images are rendered. Each image represents one eye.

3. <http://blog.irayrender.com/post/76948894710/compositing-with-light-path-expressions>

Left over right

The image of the left eye is rendered above the image of the right eye. The rendered images are twice the height of a non-stereo image.

Left next to right

The image of the left eye is rendered beside (left side of) the image of the right eye. The rendered images are twice the width of a non-stereo image.

Eye separation distance

The deviation separating the two perspectives. Use a minor deviation that matches the natural perspectives that both eyes receive in binocular vision. The value is specified in scene units.

5.7 Archiving jobs

The Done list is a list of all ended jobs. You can clear the job list in one operation by using the Clear button.

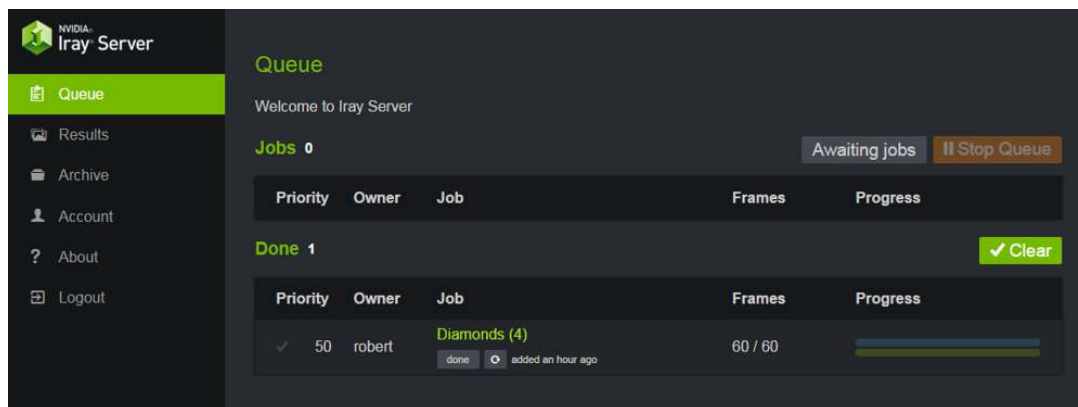


Fig. 5.3 - Clear the Done list by clicking the Clear button

Copies of the cleared jobs are stored in the archive. The render results are not deleted.

Note: You can clear and archive your own jobs only, unless you have administrator privileges.

To archive all jobs in the Done list:

1. From the Queue page beside the header Done: Click Clear. A confirmation window is displayed.
2. From the Confirmation window, click Clear jobs.
3. To view archived jobs, go to the navigation bar on the left and click Archive. The [Archive page](#) (page 41) is displayed.

6 Managing the queue

This chapter describes to manage the Iray Server queue.

6.1 Stopping a job

Before you can edit, copy, or remove a job in the queue, you must stop it. You can stop a running job as well as a job waiting in the queue. Note that a waiting job, just like a stopped job, can be edited, copied, removed. Active (running) jobs can be copied.

To stop a running or waiting job:

- From the Jobs list on the Queue page, locate the job and click **Stop** on the toolbar.

Note the following:

- When you stop a *running* job, the current state of the job is saved and the job remains in the queue. If the rendering of the current frame is not complete, it is discarded.
- When you stop a *waiting* job, Iray Server may continue to move it up the queue (depending on its priority) but it will not run the job until you restart it. In the meantime, other jobs in the queue will continue to be executed.

6.2 Restarting a job

To restart a job:

- From the Jobs list on the Queue page, locate the job and click **Start** on the toolbar.

The job is restarted from its current position in the render queue.

6.3 Moving a job in the queue

The position of a job in the queue is determined by its priority level. You change the position of a job in either of the following ways:

- From the Queue page, find the job you want to move up or down the queue and drag-and-drop the Priority icon.

If you drag and drop a job with a low priority in front of a job with a higher priority, the priority level is changed. If jobs have the same priority level, their position in the queue is determined by waiting time, unless you manually drag and drop the job.

- From the Edit page, change the priority level as needed.

When you close the Edit page and restart the job, the job order reflects your update.

Note the following:

- Zero (0) is the highest priority level supported by Iray Server
- The highest priority level you can assign to a job is defined by your administrator
- There is no minimum priority level
- If jobs have the same priority level, their position in the queue is determined by waiting time

6.4 Deleting (removing) a job in the queue

To delete (remove) a job in the queue:

1. From the Queue page, locate the job you want to delete from the queue.
2. Check the status of the job. If the job is active, you must [stop](#) (page 37) it before removing it.
3. Click Remove and confirm the operation.

7 Managing results

This chapter describes the Results page, where you can view, download, and delete job results:




Fig. 7.1 - The Results page. By default, only your rendering results are displayed. To display the rendering results of other users, you need administrator rights.

The table displayed on the Results page provides a complete list of render results, which includes:

- Completed jobs
- Uncompleted jobs where at least one frame was rendered
- Archived jobs

When searching the table, note the following:

- Jobs are listed from most recent to oldest.
- Thumbnail images identify the content of each rendered file. Use the left-right arrows to scroll through long lists of thumbnails.
- Click on a thumbnail image to display the image with a higher resolution. The file name, the type of generated output, and a Download button are displayed with this image.

The following action buttons are provided:

Delete (X)

Click this button to delete the complete results for the job. You cannot delete a selection of files.

Files

Click this button to display the results as a file list:

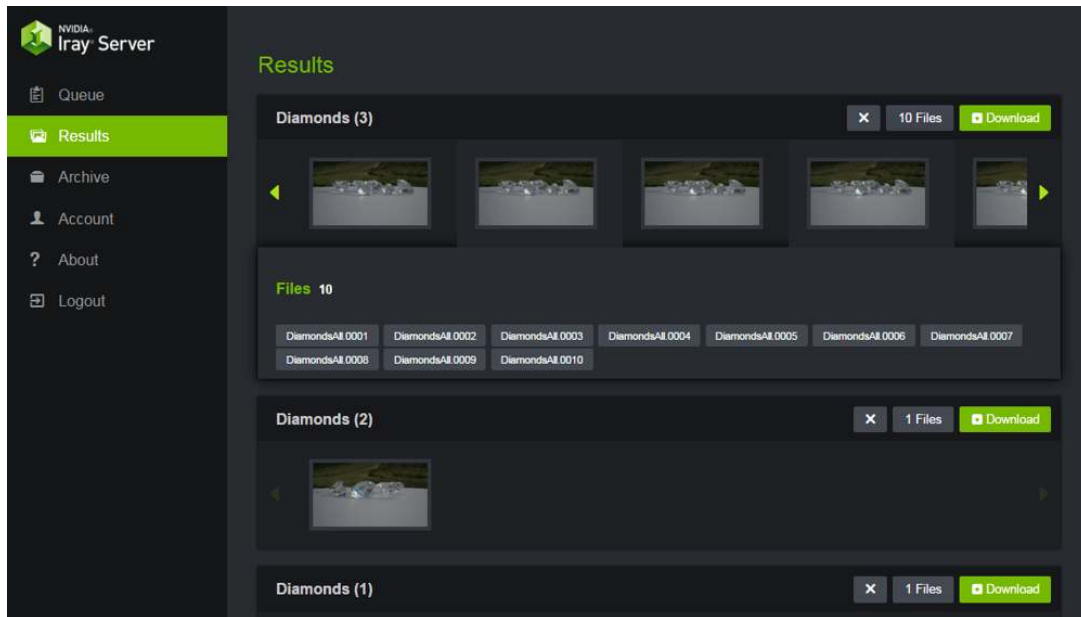


Fig. 7.2 - List of image files generated for a job. File names reflect the type of generated output

To display a particular file as full-sized image, click the file name. To download a file, select it and click **Download**.

Download

Click this button to download the complete results for a job.

A download window is displayed with the name of the zipped folder. The name of the zipped folder is the same as the job name. Specify a download location and click **OK**.

To download render files while a job is running, use the bar at the top of the Results page.

8 Managing the archive

This chapter describes the Archive page, where you can view, delete, and add archived jobs to the render queue:




Fig. 8.1 - The Archive page. By default, only your archived jobs are displayed. To display the archived jobs of other users, you need administrator rights.

The table displayed on the Archive page provides a complete list of archived results. When searching the table, note the following:

- Jobs are listed from oldest to most recent.
- The table uses the same layout as the Done table on the Queue page.

The following action buttons are provided:

Delete all

Click this button to delete all archived jobs.

Note:

- You cannot delete selected archived jobs.
- There is no mechanism to retrieve archived jobs after you delete them.

Copy

Click this button to display the [Copy](#) (page 32) window where you can edit the archived job and add it to the job queue.

Part 3 Managing settings, resources, and users

To perform the tasks described in the following chapters, you need administrator rights.

9 Managing configuration settings

This chapter describes how to specify the configuration settings for your Iray Server setup.

9.1 Touring the Settings page

You specify configuration settings from the Settings page. The Settings page is divided into the following parts:

General settings

Where you specify the server name and the message of the day

Results settings

Where you customize render results.

Events settings

Where you can add custom commands for post-processing.

9.2 Specifying general settings

From General settings you can change the default server name and the message of the day.




Fig. 9.1 - General settings

Server Name

The server name. The default name is Iray Server. If your Iray Server configuration consists of multiple clusters, it is recommended that you give the master server of each cluster a unique name to avoid any confusion.

Message of the Day

A message for all users, which is displayed near the top of the Queue page. Such messages are useful, for example, to notify users of upcoming events or server availability.

9.3 Configuring results settings

The following sections describe how to configure settings for render results and what to consider when making your choices.

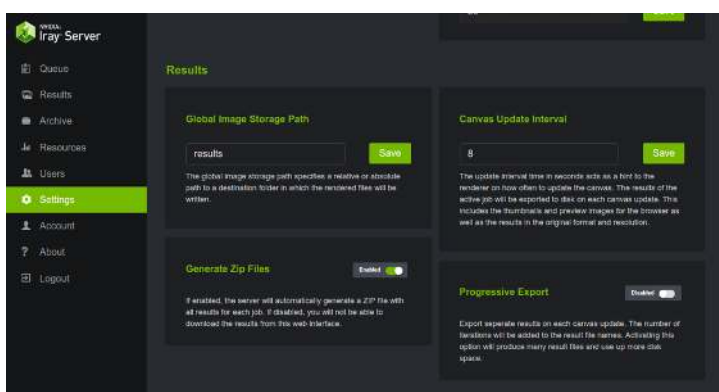


Fig. 9.2 - Resource settings

9.3.1 Changing the global storage directory

The entry field in the Global Image Storage Path box, displays the current directory path where render results are stored. The default location uses the current working directory of Iray Server, where output is stored in the results directory.

To change the global storage location:

- Type the new path in the entry field and click Save. Render results will be saved to this directory.

You can set up a global storage path or individual paths for users. Custom storage paths defined for users are unaffected by path or name changes to the global destination folder.

When specifying directory paths, you may need to specify different paths for multiple hosts. For example:

- When Iray Server runs under a Linux host and your host runs under Windows, the storage path is different, although both hosts may access the same file server.
- When using multiple Windows hosts, the setup may differ. For example, a shared storage location may be accessed as different disk drives (for example, d:\ and e:\).

On the host where Iray Server is running, the path for Iray Server is always valid. For this reason, it is recommended that you log in to that host when determining the storage path.

9.3.2 Creating zip files for downloads

You must package render results as `.zip` files before you can download them.

To generate `.zip` files:

- From the **Generate Zip Files** box, set the slider to **Enabled**. Download `.zip` files from the **Results** page.

Because the generation of `.zip` files takes time and space, consider disabling this setting when rendering locally.

9.3.3 Specifying an update interval

To specify an update interval for progressive rendering:

- From the **Canvas Update Interval** box, specify the update value in the entry field and click **Save**. The value you enter represents the number of seconds between updates.

When the image is updated, the interim result is exported to disk. Each exported result includes the results in the original format and resolution together with thumbnails and preview images.

9.3.4 Buffering interim results

To enable separate buffers for interim results of a progressive rendering:

- In the **Progressive Export** box, set the slider to **Enabled**. The results of interim renders are saved to separate buffers.

When enabled, this setting produces many result files and uses more disk space.

9.4 Adding custom commands




Fig. 9.3 - Event settings

After a job is finished, you can call custom scripts. You may use custom scripts, for example, to forward a job to a post-processing tool chain or to send notifications to users when jobs are completed. Event settings are shared by all users.

You can pass the following variables to your scripts:

<i>Variable</i>	<i>Value</i>
@user	The user who submitted the job
@job	The name of the job
@status	The status of the job, which can be either done or failed
@email	The email address of the user who submitted the job
@error	A description of the error, if any.

For Windows, use batch scripts. For Linux, use shell scripts.

Setting events is recommended for advanced users only.

10 Starting and stopping the render queue

This chapter describes how to start and stop the render queue and switch between streaming and queuing modes. These tasks are performed from the **Queue** page.

10.1 Starting the render queue

To start the render queue:

- From the Queue page, click **Start Queue**. The render queue is started and the **Stop Queue** button is displayed.

10.2 Stopping the render queue

To stop the render queue:

- From the Queue page, click **Stop Queue**. The currently running job, if any, is completed, rendering is stopped, and the **Start Queue** button is displayed.

10.3 Switching between streaming and queuing modes

You can choose between streaming and queuing modes but you cannot run both modes simultaneously. Streaming mode can accept jobs for later processing in the queue.

If a user is streaming from the server when you try to start the queue mode, a warning message is displayed:




Fig. 10.1 - Streaming warning

Do one of the following:

1. Click **Start Queue** to stop streaming mode and start queuing mode, or
2. Click **Cancel** to allow the user to continue streaming.

11 Managing rendering resources

This chapter describes how to manage rendering resources. These tasks are performed from the Resources page.

11.1 Touring the Resources page

The Resources page provides the following reports about Iray Server resources to enable you to debug and manage these resources for the optimal performance of rendering jobs:

Supported Versions

Supported versions of [Iray Bridge](#) (page 51)

Cache Manager

Configuration options for the [Cache Manager](#) (page 52)

Storage

Available storage for [render results](#) (page 52)

Hosts

A list of [servers](#) (page 53) in your Iray Server cluster and their resource status.

Logs

A downloadable [log](#) (page 53) of the master node in your Iray Server cluster

Benchmarks

[Rendering jobs](#) (page 54) for testing your configuration and its performance

11.2 Checking the supported versions of Iray Bridge

The Supported Versions report in Figure 11.1 is derived from the log files. It provides details about the supported versions of Iray Bridge for debugging and resource management tasks.

ID	Product	Hostname	IP	Cluster Interface	Admin Page	State
8	Iray trunk (310029)	trayg	172.16.1.13	0.0.0.0:10016	-	Rendering
10	Iray trunk (310029)	trayh	172.16.1.15	0.0.0.0:10020	-	Rendering
15	Iray trunk (310029)	traya	172.16.0.240	0.0.0.0:10030	-	Rendering
17	Iray trunk (310029)	traye	172.16.1.9	0.0.0.0:10034	-	Rendering
25	Iray trunk (310029)	trayb	172.16.1.3	0.0.0.0:10050	-	Rendering

Fig. 11.1 - The Supported Versions report

Iray Server launches an Iray Worker for each version of Iray that is supported. When a request is submitted to Iray Server, the job is automatically sent to the correct Iray Worker for processing. The Bridge protocol is used to associate it to the relevant version of Iray being run by an

Iray Worker. These Bridge protocols, which are referenced in the log file, may help you track down multi-version issues (if any).

For each supported version of Iray Bridge, the report provides following information:

ID

The unique identifier for each worker

Product

The installed version of Iray

Host name

The host on which the worker is installed

IP

The host IP address

Cluster interface

The IP and port number used by Iray on the host

Admin page

The value generated when the `--admin-server` command is used to start Iray Server

State

The current rendering state of the worker, which can be either `rendering` or `idle`

11.3 Checking the cache manager settings

The Cache Manager report identifies the current settings for the Cache Manager options:

Address

The address to listen on. The only time you should change the address is when there is a conflict in the port usage.

Data path

The directory used by the cache manager to store its data.

Low water mark

Garbage collection ends when memory usage (in MB) falls below this value.

High water mark

Garbage collection starts when memory usage (in MB) exceeds this value. Garbage collection is disabled when this value is set to zero (0).

11.4 Checking the available storage for render results

The Storage report specifies the available disk space for render results. For each storage location, it provides the following information:

Total space

The amount of space on the hard disk

Available space

The amount of space that is still free

Path

The directory path where render results are stored for all users, unless a custom path is specified

11.5 Checking the server configuration

The Hosts report in Figure 11.2 lists the servers in your Iray Server configuration. The configuration mode, protocol, and IP address are displayed at the top-right of the report. To enable and disable CPU and GPU usage, click on the host name, and select the appropriate choice.

ID	Hostname	Memory	CPU Cores	GPUs	Uptime
1	master trayh	24 GB	16	4	3 min
2	trayg	24 GB	16	4	3 min
3	traya	47 GB	16	2	2 min
4	traye	20 GB	16	2	1 min
5	trayb	47 GB	16	2	1 min

Fig. 11.2 - The Hosts report, which also displays the configuration mode (top-right)

For each host, the following information is provided:

- ID
- Host name
- Memory
- CPU cores
- GPUs
- Uptime


11.6 Checking and downloading the log of the master node

The log of the master node is displayed in a scrollable viewport. Note the following:

- The most recent log messages are displayed at the bottom of the window.
- The log is locked by default. When the log is locked, scrolling is automatic. For manual scrolling, you need to unlock the scroll bar by clicking the lock icon.
- To download the log, click Download.

11.7 Benchmarks

If the installation directory contains a folder with benchmark scenes, a table with the scene names is displayed:



The screenshot shows a dark-themed interface with a title 'Benchmarks' in green. Below the title is a table with two columns: 'Name' and '+ Add to Queue'. The table lists 12 benchmark scenes, each with a corresponding button to add it to the queue.

Name	+ Add to Queue
cubes	+ Add to Queue
5_spheres	+ Add to Queue
animals	+ Add to Queue
buddha	+ Add to Queue
bunny-backplate	+ Add to Queue
cornell_box	+ Add to Queue
decals	+ Add to Queue
interior	+ Add to Queue
light_glow	+ Add to Queue
matroschka	+ Add to Queue
shader_ball	+ Add to Queue
vases_horse	+ Add to Queue

Fig. 11.3 - Benchmark scenes

To run a benchmark scene, click **Add to Queue**. The scene is added as a new job to the queue.

A benchmark job uses the following predefined rendering options:

- Format: PNG
- Resolution: 1920x1080
- Maximum iterations: 1000

You can specify which version of Iray you want to use for rendering the job.

12 Managing users

This chapter describes how to manage user access to Iray Server. These tasks are performed from the Users page.

12.1 Touring the Users page

The Users page provides a complete and editable list of users and user settings:




Fig. 12.1 - The Users page

For each user, the following information is displayed:

Name

User ID.

Admin

Administrator privileges: Yes or no.

Max. Priority

The highest priority level that a user can assign to a job in the queue. The highest priority level is 0 (zero). The lowest priority level is 100.

Image Storage

Specifies the directory where render results are stored for this user.

Action

Specifies the available editing actions for each user.


You can organize user accounts in the following ways, where:

- A unique account name and password is defined for each user
- A project account is defined and account users share a password

12.2 Adding a user

To add a user:

1. Click the **Add user** button near the top of the page. The New User window is displayed:



New User

Name
A-Z a-z 0-9 - _ @

E-mail

Password
At least 5 characters

Max. priority
Max. job priority - lower values mean higher priority

Permissions
Normal user or administrator

Require password change
Prompt the user to change the password on next login

Image storage path
Absolute or relative path to a folder to store the results of the new user

User name

E-mail address

Password

Confirm Password

100

Normal user

Yes

Destination path for results

Cancel Save

Fig. 12.2 - The New User window

2. Fill in the appropriate fields.
3. Click **Save**. A confirmation message is displayed in the top-right corner of the page. If the Save operation is successful, the new user is added to the Users list.

12.3 Deleting a user


To delete a user:

1. In the Users list locate the user.
2. In the Action column, click **Delete**. A confirmation window is displayed.
3. Click **Remove** to delete the user.

12.4 Editing a user profile

To edit a user profile:

1. In the Users list locate the user.
2. In the Action column, click **Edit**. The Edit User window is displayed:



Edit User

Name

E-Mail

Max. Priority

Permissions Admin

Image Storage Path Custom

Fig. 12.3 - The Edit User window

3. After you finish your edits, click **Save**.

13 Troubleshooting

This chapter describes how to resolve errors that may arise when running Iray Server.

13.1 Clustering

Q: I am using UDP with multicast and it is not working correctly.

A: Iray Server supports two cluster modes: UDP with multicast and TCP/IP discovery mode. UDP with multicast is generally quite easy to set up. You simply enable it; you do not need to concern yourself with addresses.

Some networks, however, either do not support UDP with multicast or support very low bandwidth only. In such cases, TCP/IP discovery mode is recommended. For more information about cluster modes, refer to [Cluster options](#) (page 20).

Q: The cluster isn't forming when a host has more than one connected interface.

A: You should use the `--interface` to explicitly specify the network interface that you want to use on the hosts that have more than one connected interface. For more details about the option the [Cluster options](#) (page 20).

Q: Workers will not start and the following error message is displayed:

```
Error: connection timeout for worker
"C:\Program Files\NVIDIA Corporation\Iray Server\workers\XXXXXXX.XXXX"
```

A: Antivirus software can misinterpret websocket communication as a thread. Ensure Iray Server installation paths are listed as exemptions. For guidance, refer to behavior monitoring and scan exclusion in your antivirus software documentation.

13.2 Farming

Q: A job has ended before finishing the rendering operation.

A: If a worker crashes during rendering, the whole job fails; the master node does not ask the affected render node to restart the worker.

Restart the job with the Queue. Only the frames that failed will be rendered.

13.3 Out-of-memory

Q: A job appears to hang which prevents Iray Server from uploading the job correctly.

A: The computer on which Iray Server is running may be low on memory. Check the log for typical error messages such as `error: Socket write operation failed.` and `Failed to read data for element ...`. Under Windows, select the Performance tab in the Task Manager window. Verify memory usage.

If memory usage is extremely high, either shutdown unnecessary applications and services or run Iray Server on another host.

