

Qube Movie Job Submission

Job Submission Details



Not all sections need to be filled in in order to render. Only the fields marked in red are required.

Qube Job Basics	
Name	<input type="text"/>
Priority	9999
Instances	1
Max Instances	-1

[Click here for details...](#)

Name

This is the name of the job of the job so it can be easily identified in the Qube! UI.

Priority

Every job in Qube is assigned a numeric priority. Priority 1 is higher than priority 100. This is similar to 1st place, 2nd place, 3rd place, etc. The default priority assigned to a job is 9999.

Instances

This is the number of copies of the application that will run at the same time across the network. The combination of "Instances=1" and "Max Instances=-1" means that this job will take as much of the farm as it can, and all jobs will share evenly across the farm.

Examples:

On a 12 slot(core) machine running Maya if you set "Instances" to 4

"Reservations" to "host.processors=3"

Qube! will open 4 sessions of Maya on the Worker(s) simultaneously, which may consume all slots/cores on a given Worker.

if you set

"Instances" to 1

"Reservations" to "host.processors=1+"

Qube will open 1 session of Maya on a Worker, consuming all slots/cores ("host.processors=1+" is used for all slots/cores).

Max Instances

If resources are available, Qube! will spawn more than 'Instances' copies of the application, but no more than 'Max Instances'. The default of -1 means there is no maximum. If this is set to 0, then it won't spawn more than 'Instances' copies.

More on [Instances & Reservations & SmartShare Studio Defaults](#)

Required parameters	
exe	<input type="text" value="images_to_movie.py"/> <input type="button" value="Browse"/>
jobid	<input type="text" value="0"/>
transcoder	<input type="text"/>
transcoder_exe	<input type="text"/> <input type="button" value="Browse"/>
transcoder_command	<input type="text"/>
transcoder_frameSyntax	<input type="text"/>
movie_ext	<input type="text" value="mov"/>
movie_path	<input type="text"/>
image_filter	<input type="text"/>

[Click here for details...](#)

exe

Explicit path to executable that will generate the movie file(s). **Important:** Be aware that if you are submitting from one OS to a different one, the path will be different.

jobid

The Qube! jobid that will create the images to be converted to a movie. If this window is generated from the "Generate Movie" option of an application submission UI, then this section will be prefilled.

transcoder

Available transcoders. Choose from the list of available transcoders. This dropdown can be edited through the qube_imagesToMovie.py file in the simpleCMDs directory

transcoder_exe

The path to transcoder executable. **Important:** Be aware that if you are submitting from one OS to a different one, the path will be different.

transcoder_command

The transcoder command to run, including any command-line arguments and flags.

transcoder_frameSyntax

The syntax used to indicate frames for the transcoder (ie. # or %4d). If this window is generated from the "Generate Movie" option of an application submission UI, then this section will be prefilled.

movie_ext

The extension of the resulting movie file (e.g., MOV). **Important:** When choosing the output format be aware that extensions such as .MOV or .AVI can not be distributed across the farm. That means that this job should be run only on one Worker as these file formats can not be split across the farm.

movie_path

The default output location is the to the same directory as the images the movie was created from. This field allows you to override that. You can enter an output filename, or a directory (which is indicated by ending the name with a "/"). You can indicate locations relative to the input image directory, e.g., "../"

image_filter

A regular expression (regex) for choosing the images to put into the movie. This is useful if the image directory contains multiple different layers, and you want to select just one set. For example, "final*.jpg" would select only the images in that directory that begin with the word "final" and end with ".jpg"

Qube SimpleCmd and Shell Parameters

Cmd Template

```
%(executable)s -project "%(Project_path)s" -s  
QB_FRAME_START -e QB_FRAME_END -i  
QB_FRAME_STEP
```

Shell (Linux/OSX)

[Click here for details...](#)

Cmd Template

This is used to create the command string for launching the job on the worker. It will be set differently depending on the application you are launching from.

Shell (Linux/OSX)

Explicitly specify the Linux/OS X shell to use when executing the command (defaults to /bin/sh).

Qube Job Tags

Job Show

Job Shot

Job Seq

Job Client

Job Dept

Job Custom1

Job Custom2

Job Custom3

Job Custom4

Job Custom5

[Click here for details...](#)

Qube Job Tags

New in Qube 6.5

Note: The Job Tags section of the submission UI will not be visible unless they are turned on in the [Preferences](#) in the Wrangler View UI. Job Tags are explained in detail on the [Job Tags](#) page.

Qube Worker Selection	
Hosts	<input type="text"/> <input type="button" value="Browse"/>
Groups	<input type="text"/> <input type="button" value="Browse"/>
Omit Hosts	<input type="text"/> <input type="button" value="Browse"/>
Omit Groups	<input type="text"/> <input type="button" value="Browse"/>
Priority Cluster	<input type="text"/> <input type="button" value="Browse"/>
Host Order	+host.processors.avail <input type="button" value="Browse"/>
Requirements	<input type="text"/> <input type="button" value="Browse"/>
Reservations	<input type="text"/> <input type="button" value="Browse"/>
Restrictions	<input type="text"/> <input type="button" value="Browse"/>

[Click here for details...](#)

Hosts

Explicit list of Worker hostnames that will be allowed to run the job (comma-separated).

Groups

Explicit list of Worker groups that will be allowed to run the job (comma-separated). Groups identify machines through some attribute they have, eg, a GPU, an amount of memory, a license to run a particular application, etc. Jobs cannot migrate from one group to another. See [worker_groups](#).

Omit Hosts

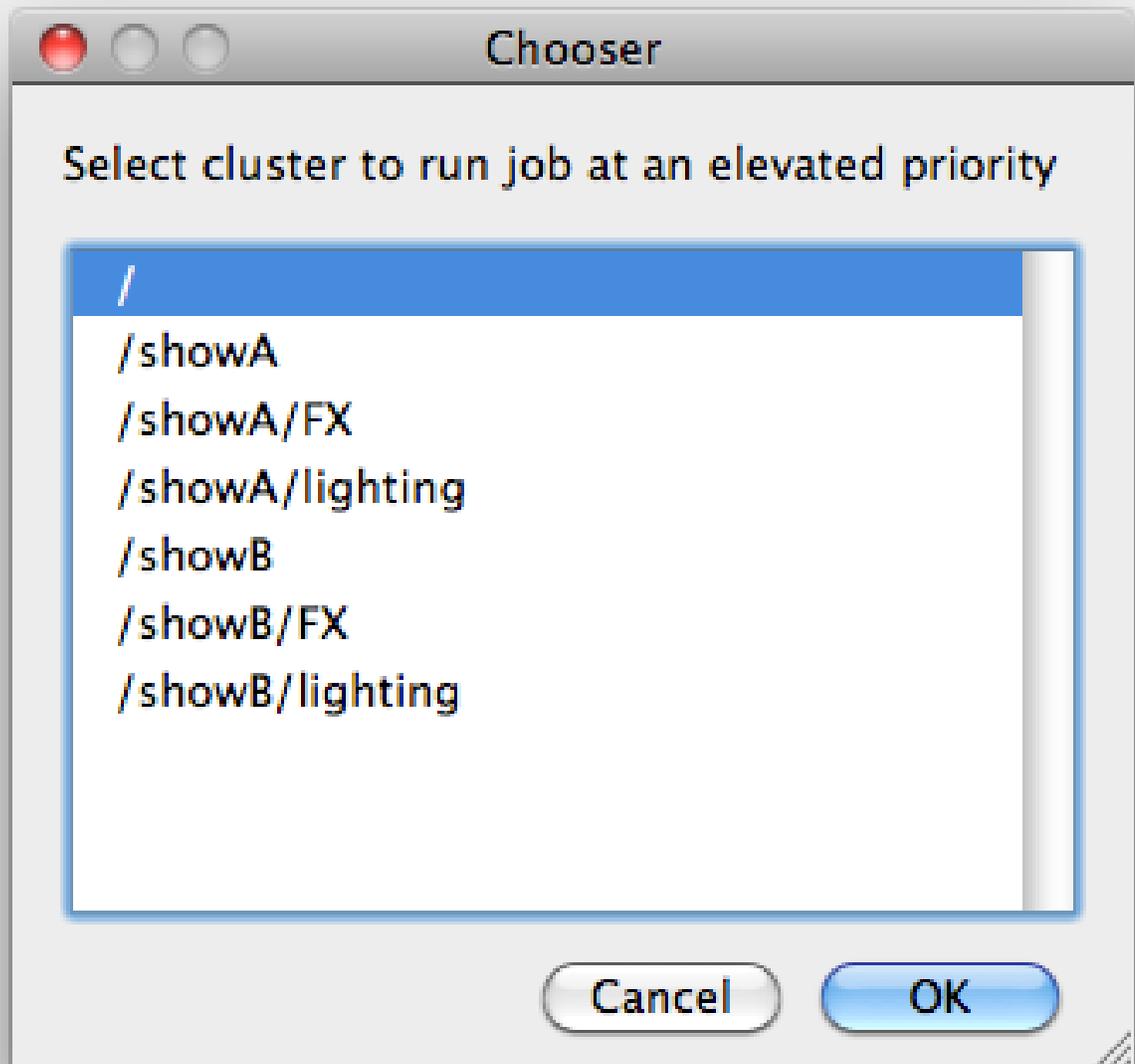
Explicit list of Worker hostnames that are **not** allowed run the job (comma-separated).

Omit Groups

Explicit list of Worker groups that are **not** allowed to run the job (comma-separated).

Priority Cluster

Clusters are non-overlapping sets of machines. Your job will run at the given priority in the given cluster. If that cluster is full, the job can run in a different cluster, but at lower priority. [Clustering](#)



Example:

- A job submitted to /showB/lighting will run with its given priority in /showB/lighting cluster.
- If /showB/lighting is full, that job can run in /showB/FX, but at a lower priority.
- If both /showB/lighting and /showB/FX are full, the job can run in /showA/* at an even lower priority.

Host Order

Order to select Workers for running the job (comma-separated) [+ means ascending, - means descending].

Host Order is a way of telling the job how to select/order workers

- "+host.processors.avail" means prefer workers which have more slots available
- "+host.memory.avail" means prefer workers which have more memory available
- "+host.memory.total" means prefer workers which have more total memory
- "+host.processor_speed" means prefer workers with higher cpu speeds
- "+host.cpus" means prefer workers with higher total cpu slots

Requirements

Worker properties needed to be met for job to run on that Worker (comma-separated, expression-based). Click 'Browse' to choose from a list of Host Order Options.

Requirements is a way to tell the workers that this job needs specific properties to be present in order to run. The drop-down menu allows a choice of OS:

- "winnt" will fill the field with "host.os=winnt" which means only run on Windows based workers
- "linux" will fill the field with "host.os=linux" which means only run on Linux based workers
- "osx" will fill the field with "host.os=osx" which means only run on OSX based workers

You can also add any other Worker properties via plain text. Some examples:

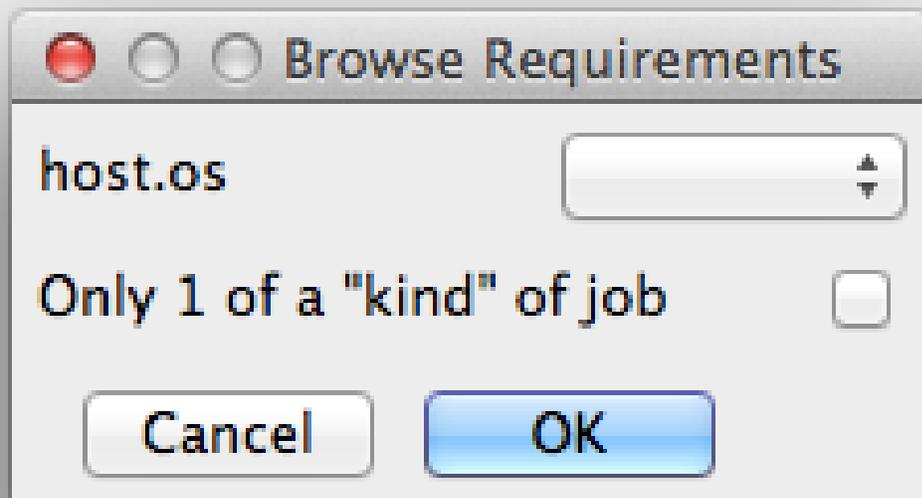
- "host.processors.avail.=4" means only run this job on workers that have 4 or more slots available
- "host.processors.used=0" means only run this job on workers with 0 slots in use
- "host.memory.avail=400" means only run this job on workers that have 400 memory available

With integer values, you can use any numerical relationships, e.g. =, <, >, <=, >=. This won't work for string values or floating point values. Multiple requirements can also be combined with AND and OR (the symbols && and || will also work).

The 'Only 1 of a "kind" of job' checkbox will restrict a Worker to running only one instance with a matching "kind" field (see below). The prime example is After Effects, which will only allow a single instance of AE on a machine. Using this checkbox and the "Kind" field, you can restrict a Worker to only one running copy of After Effects, while still leaving the Worker's other slots available for other "kinds" of jobs.

Reservations

Worker resources to reserve when running job (comma-separated, expression-based).



Browse Reservations

host.memory	<input type="text" value="0"/>	<input type="button" value="↑"/> <input type="button" value="↓"/>	Mb
host.processors	<input type="text" value="1"/>	<input type="button" value="↑"/> <input type="button" value="↓"/>	<input type="checkbox"/> All
global.vray	<input type="text" value="0"/>	<input type="button" value="↑"/> <input type="button" value="↓"/>	
global_host.nuke	<input type="text" value="0"/>	<input type="button" value="↑"/> <input type="button" value="↓"/>	
license.prman	<input type="text" value="0"/>	<input type="button" value="↑"/> <input type="button" value="↓"/>	

OK Clear Cancel

Reservations is a way to tell the workers that this job will reserve the specific resources for this job.

Menu items:

- "host.processors" this will fill the field with "host.processors=X" which means reserve X slots on the worker while running this job
- "host.memory" this will fill the field with "host.memory=X" which means only reserve X memory on the worker while running this job

Other options:

- "host.license.nuke=1" when a [Global Resources](#) entry has been made you can reserve any arbitrary named item. **New in 6.6:** Once you global resource, it will show up in this menu (eg global.vray above).
- See also [Job Reservations](#)

Restrictions

Restrict job to run only on specified clusters ("|" -separated) [+]

means all below, * means at that level]. Click 'Browse' to choose from a list of Restrictions Options.

Restrictions is a way to tell the workers that this job can only run on specific clusters. You can choose more than one cluster in the list.

Examples:

- Choosing /showA would restrict the job to machines that are only in the /showA cluster, and no other cluster, not even those below /showA.
- Choosing /showA/* would restrict the job to the cluster(s) *below* /showA, but *not including* /showA
- Choosing /showA/+ would restrict the job to /showA and all the clusters below it.

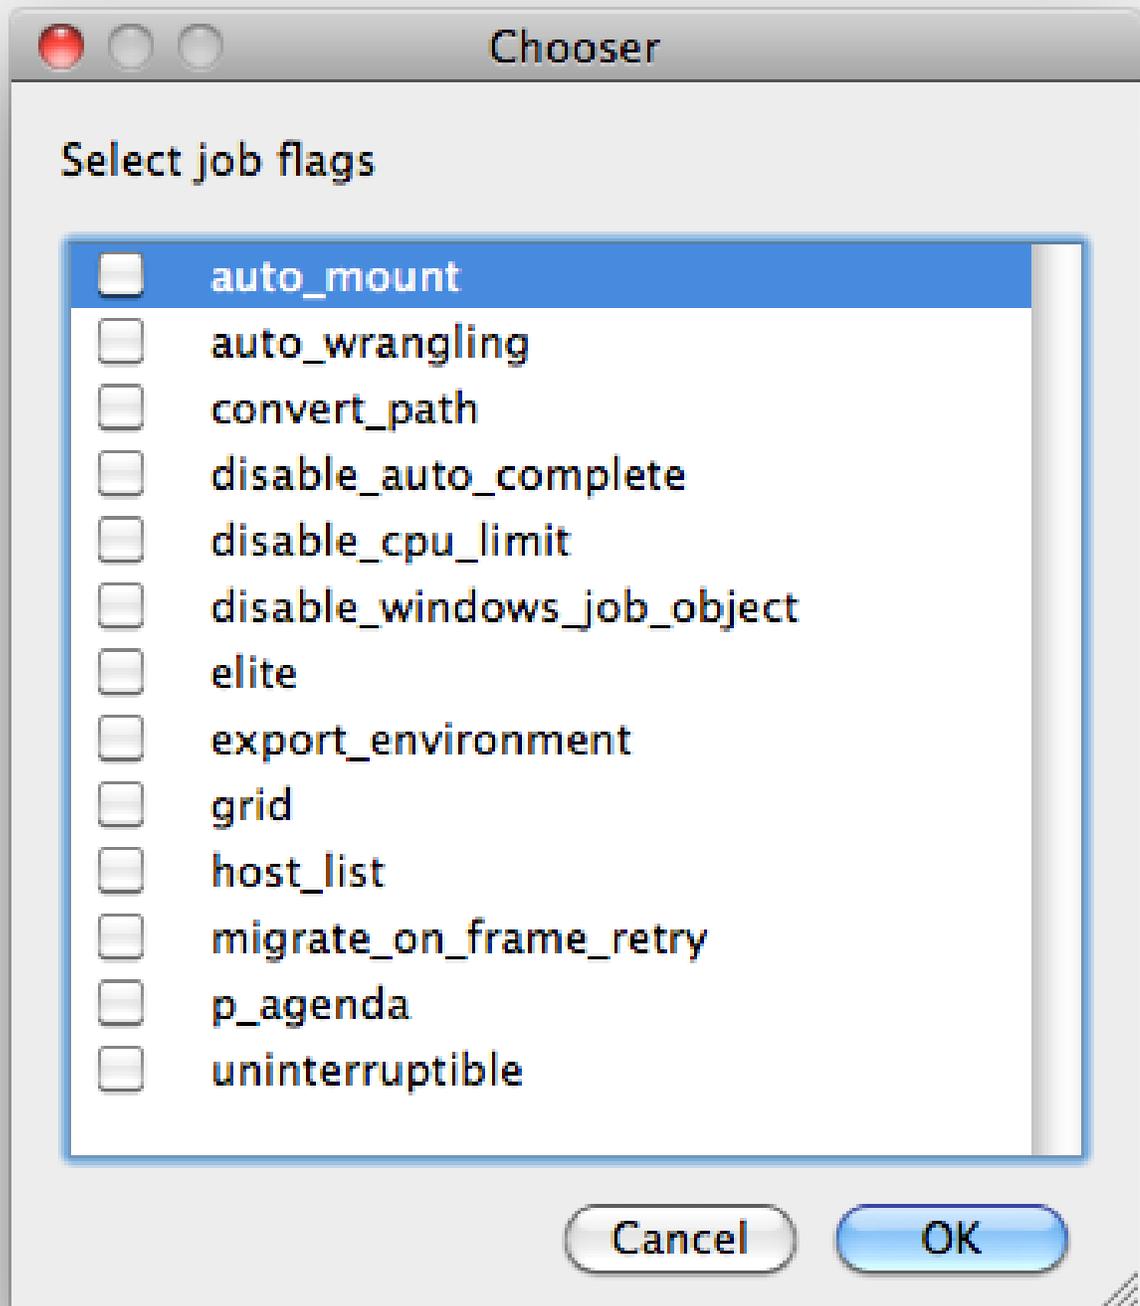
See Also

- [Controlling Host Selection](#)
- [How to use qbrk.conf](#)
- [worker_groups](#)
- [worker_cluster](#)
- [How to use clustering for workers](#)

[Click here for details...](#)

Flags

List of submission flag strings (comma separated). Click 'Browse' to choose required job flags.

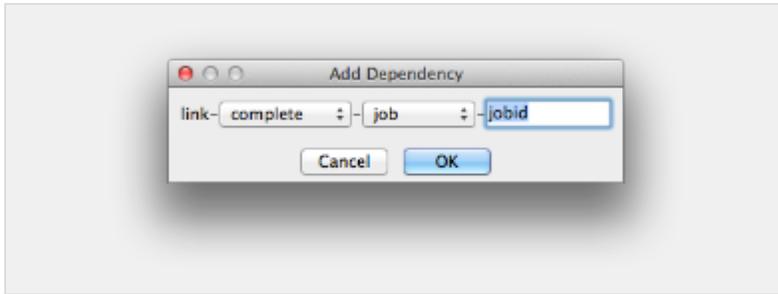


See [this page](#) for a full explanation of flag meanings

Dependency

Wait for specified jobs to complete before starting this job

(comma-separated). Click 'Add' to create dependent jobs.



You can link jobs to each other in several ways:

- "complete" means only start this job after designated job completes
- "failed" means only start this job if the designated job fails
- "killed" means only start this job if the designated job has been killed
- "done" means start this job if the designated job is killed/failed/complete

The second menu chooses between "job" (the entire set of frames) and "work" (typically a frame). So to link frame 1 of one job to frame 1 of a second, job, you would choose "work" in this menu. If you want to wait for all the frames of one job to complete before starting a second, then choose "job". The other option, "subjob", refers to the instance of a job. This is much less common, but means that, for example, the instance of Maya that was running frames has completed.

For a complete description on how to define complex dependencies between jobs or frames, please refer to the [Callbacks](#) section of the Developers Guide.

Email (job complete)

Send email on job completion (success or failure). Sends mail to the designated user.

Email (failed frames)

Sends mail to the designated user if frames fail.

Blocked

Set initial state of job to "blocked".

Stderr->Stdout

Redirect and consolidate the job stderr stream to the stdout stream. Enable this if you would like to combine your logs into one stream.

Job Label

Optional label to identify the job. Must be unique within a Job Process Group. This is most useful for submitting sets of dependent jobs, where you don't know in advance the job IDs to depend on, but you do know the labels.

Job Kind

Arbitrary typing information that can be used to identify the job. It is commonly used to make sure only one of this "kind" of job runs on a worker at the same time by setting the job's requirements to

include "not (job.kind in host.duty.kind)". See [How to restrict a host to only one instance of a given kind of job, but still allow other jobs](#)

Process Group

Job Process Group for logically organizing dependent jobs. Defaults to the jobid. Combination of "label" and "Process Group" must be unique for a job. See [Process group labels](#)

Retry Frame/Instance

Number of times to retry a failed frame/job instance. The default value of -1 means don't retry.

Retry Work Delay

Number of seconds between retries.

Subjob Timeout

Kill the subjob process if running for the specified time (in seconds). Value of -1 means disabled. Use this if the acceptable instance/subjob spawn time is known.

Frame Timeout

Kill the agenda/frame if running for the specified time (in seconds). Value of -1 means disabled. Use this if you know how long frames should take, so that you can automatically kill those running long.

The screenshot shows a window titled "Qube Job Environment". It contains three main sections: "Cwd" with a text input field, "Environment Variables" with a table, and "Impersonate User" with a text input field. The table has two columns: "Key" and "Value".

Key	Value

[Click here for details...](#)

Cwd

Current Working Directory to use when running the job.

Environment Variables

Environment variables override when running a job. You can specify key/value pairs of environment variables

This is useful when you might need different settings for your render applications based on different departments or projects.

Impersonate User

You can specify which user you would like to submit the job as.

The default is the current user. The format is simply <username>. This is useful for troubleshooting a job that may fail if sent from a specific user.

Example:

Setting "josh" would attempt to submit the job as the user "josh" regardless of your current user ID.

Note: In order to do this, the submitting user must have "impersonate user" permissions.

Key	Value

Key	Value

Key	Value

[Click here for details...](#)

Windows-only Environment Variables

Used to provide OS specific environment variables for Windows. Enter variables and values to override when running jobs.

Linux-only Environment Variables

Used to provide OS specific environment variables for Linux. Enter variables and values to override when running jobs.

Darwin-only Environment Variables

Used to provide OS specific environment variables for OS X. Enter variables and values to override when running jobs.

Min File Size: 0

regex_highlights

regex_errors

regex_outputPaths

regex_progress

regex_maxLines: 20

[Click here for details...](#)

Min File Size

Used to test the created output file to ensure that it is at least the minimum size specified. Put in the minimum size for output files, in bytes. A zero (0) disables the test.

regex_highlights

Used to add highlights into logs. Enter a regular expression that, if matched, will be highlighted in the information messages from stdout/stderr.

regex_errors

Used to catch errors that show up in stdout/stderr. For example, if you list "error: 2145" here and this string is present in the logs, the job will be marked as failed. This field comes pre-populated with expressions based on the application you are submitting from. You can add more to the list, one entry per line.

regex_outputPaths

Regular expression for identifying outputPaths of images from stdout/stderr. This is useful for returning information to the Qube GUI so that the "Browse Output" right-mouse menu works.

regex_progress

Regular expression for identifying in-frame/chunk progress from stdout/stderr. Used to identify strings for relaying the progress of frames.

regex_maxlines

Maximum number of lines to store for regex matched patterns for stdout/stderr. Used to truncate the size of log files.



Examples

To see examples of regular expressions for these contexts, look at the Nuke (cmdline) submission dialog - it has several already filled in.

Qube Actions ▾
generateMovie

▾ [Click here for details...](#)

GenerateMovie

Select this option to create a secondary job that will wait for the render to complete then combine the output files into a movie.

Note: For this to work correctly the "Qube (ImagesToMovie) Job..." has to be setup to use your studios transcoding application.

Qube Notes ▾
Account

Notes

▾ [Click here for details...](#)

Account

Arbitrary accounting or project data (user-specified). This can be used for creating tags for your job.

You can add entries by typing in the drop-down window or select already created accounts from the drop-down.

See also [Qube! Job Tags](#)

Notes

Freeform text for making notes on this job. Add text about the job for future reference. Viewable in the Qube UI.