

# Glossary

## Entities

Tractor has 5 key entities which are tracked in its database and available for inspection in the Dashboard, the tq command-line utility, and Python API.

### Job

A *job* is a logical body of work that is submitted to the system for processing. It describes what needs to get done and in which order. It can be build to be arbitrarily large and complex, from rendering a single layer of a single frame to multiple layers over a range of frames. Its can comprise of rendering, compositing, administration, and any other process that can run on a host.

### Task

A *task* is a unit of work which partially completes a job. A task can run in parallel or in series with other tasks. Tasks and their order of operation constitute the struture of a job.

### Command

A *command* is a single program that is executed on a *blade*. It can be a render, a notification script, an administration command, or any other program that can execute on the host.

### Invocation

An *invocation* is a set of metrics that describe the execution of a command on a blade. When a command runs multiple times, it will have multiple invocations to show when it ran, where it ran, and what memory and CPU resources were used.

### Blade

A *blade* is a remote execution server, a program that runs *commands* and reports the results of the command to the engine.

## Attributes

### Job Attributes

name	type	description
jid	integer	The Job Identifier, the unique, numerical identifier of a job.
owner	string	The user that submitted the job.
spoolhost	string	The host from which the job was submitted.
spoolfile	string	The path to filename of the spooled jobfile.
spoolcwd	string	The working directory of command that spooled the job.
spooladdress	ip address	The IP address of the host that spooled the job.
title	string	The title of the job.
assignments	encoded string	These are global job variable assignments, defined in the <i>Assign</i> section of a job file.
dirmap	json structure	A map for translating paths according to architecture.
tier	string	The name of the tier. A tier is an ordered partitioning of the queue.
priority	float	The priority determines placement of the job in the queue.
crews	string list	The list of crews of the job. This will restrict the job to run on blades which are configured to accept jobs running under those crews.
projects	string list	a list of designations which affects how the active tasks are counted in sharing limits

tags	string list	The names of the limits that will govern all commands of the job. A limit will cap the number of concurrent executions of commands tagged with the limit name.
service	string	The service key expression that will govern all commands of the job. The service key expression is a boolean expression that allows commands to be matched with blades that have the capabilities or designation to run them.
envkey	string list	A list of environment key names that establish the environment in which all commands of the job will run.
edit policy	string	The name of the policy affecting which users can manipulate job.
minslots	integer	The minimum number of slots required to run a command of the job.
max slots	integer	The maximum number of slots required to run a command of the job.
eta level	integer	The level of job graph used to estimate remaining time to completion (unused).
after jids	string	The list of ids of jobs that must finish before this job is started.
max active	integer	The maximum number of concurrent active commands allowed for this job.
serial subtasks	boolean	A boolean value indicating whether subtasks are to be executed serially.
spooltime	timestamp	The time the job was spooled.
pause time	timestamp	The time job was paused.
after time	timestamp	The time after which job will be considered for scheduling.
start time	timestamp	The time job first had an active task.
stop time	timestamp	The last time the job had a task that completed.
delete time	timestamp	The time the job was deleted.
elapsed seconds	integer	The total elapsed seconds the job has been active.
estimated seconds	integer	The estimated total elapsed task seconds for the job to complete.
num tasks	integer	The number of tasks in the job.
num blocked	integer	The number of blocked tasks in the job. A blocked task requires other tasks to finish before it can become ready, and later, active.
num ready	integer	The number of ready tasks in this job. A ready task is a waiting task that does not need any other tasks to finish in order to become active.
num active	integer	The number of active tasks in the job.
num error	integer	The number of error tasks in the job.
num done	integer	The number of done tasks in this job.

max tid	integer	The highest task id of all tasks of the job, including detached ones. This is used by the engine to manage task id assignment to newly created tasks from expand tasks. Detached tasks are tasks that are not considered for scheduling because they were produced by expand tasks and the job has since been restarted.
max cid	integer	The highest command id of all commands of the job, including ones of detached tasks. This is used by the engine to manage command id assignment to newly created commands from expand tasks.
com ment	string	A user-defined job comment string.
met adata	string	User-defined metadata of the job.
pil	integer	Place In Line: ordering of job relative to other jobs of same priority; initially set to the jid.
lastn oteid	integer	id of most recent note of job

## Task Attributes

name	type	description
jid	integer	The unique identifier of the job the task belongs to.
tid	integer	The Task Identifier, the unique identifier for the task within the job.
title	string	The title of the task.
id	string	The unique string id for the task within the job.
service	string	The service key expression that will govern all commands of the task. The service key expression is a boolean expression that allows commands to be matched with blades that have the capabilities or designation to run them.
minslots	integer	The minimum number of slots required to run a command of the task.
maxslots	integer	The maximum number of slots required to run a command of the task.
cids	integer list	A list of command ids of the commands of the task.
serials ubtasks	boolean	A boolean value indicating whether subtasks are to be executed serially.
ptids	integer list	A list of task ids of the parent tasks. Only a task with associated Instances has multiple parents. All child tasks must successfully finish before their parent can become active.
attached	boolean	A boolean value which, if false, indicates the task was result of an expand task that was retried.
state	string	The task state. It can be blocked, ready, active, error, or done.
stateti me	time stamp	The time the task became its current state.
readyti me	time stamp	The time the task became ready.
activeti me	time stamp	The time that the task became active.
currcid	integer	Command id of the command that will run next, or is currently running.
haslog	boolean	A boolean indicating whether task has output in log.
preview	string list	A list of command arguments (argv) of the preview command.
chaser	string list	A list of command arguments (argv) of the chaser command.

progress	float	The task progress, expressed as a percentage value between 0 and 100. This is set when tasks emit ALF_PROGRESS messages.
metadata	string	User-defined metadata of the task.
resumeblock	boolean	A boolean indicating whether the task will automatically resume any resumable ancestor tasks.
retrycount	integer	A counter of the number of passes the task will incur if the job is left to run to completion. A task retry or job restart will increment it if the task had become active at some point.

## Command Attributes

name	type	description
jid	integer	The unique identifier of the job the command belongs to.
tid	integer	The unique identifier of the task the command belongs to.
cid	integer	The Command Identifier, the unique identifier for the command within the job.
argv	string list	The list of command arguments representing command.
local	boolean	A boolean that is true if command is to be run on spooling host.
expand	boolean	A boolean that is true if the output of command emits script defining more tasks.
runtype	string	The type of command. Possible values include "normal" and "cleanup".
msg	string	A string value which a blade emits to a pipe to the stdin of the command.
service	string	The service key expression that will govern all commands of the task. The service key expression is a boolean expression that allows commands to be matched with blades that have the capabilities or designation to run them.
tags	string list	The names of the limits that will govern all commands of the job. A limit will cap the number of concurrent executions of commands tagged with the limit name.
id	string	The unique string id for the task within the job.
refersto	string	An id of another task or command. Setting this causes the command to run on the same blade that ran the referred task or command ran on.
minslots	integer	The minimum number of slots required to run the command.
maxslots	integer	The maximum number of slots required to run the command.
envkey	string list	A list of environment key names that establish the environment in which the command will run.
retrycodes	integer list	A list of return codes that will trigger auto-retry logic for the command.
metadata	string	User-defined metadata for the command.
resumewhile	string list	A command argument list that is executed by a blade, or a list of special keywords; either of which determine whether the command is resumable.
resumeonpin	boolean	A boolean value indicating whether the command should run on the same host when it is resumed.
minrunsecs	float	The minimum number of second for command to run to not be considered an error.
maxrunsecs	float	The maximum number of second for command to run before being killed.

## Invocation Attributes

name	type	description
jid	integer	The unique identifier of the job the invocation belongs to.
tid	integer	The unique identifier of the task the invocation belongs to.
cid	integer	The unique identifier of the command the invocation belongs to.
iid	integer	The Invocation Identifier, the unique identifier for the invocation within the command.
current	boolean	A boolean value that is true if this is the most recent invocation. When a task is retried, no existing invocations will be considered current.
numslots	integer	The number of slots used by the invocation.
limits	string list	A list of limits in use by the invocation. This value is used to reconstruct the limit counters from currently running invocations should the engine be restarted.
starttime	timestamp	The start time of the invocation.
stoptime	timestamp	The stop time of the invocation.
pid	integer	The process id of the invocation.
rss	float	The resident set size of the process, in GB.
mem	float	The memory usage of the process, in GB.
cpu	float	The current CPU utilization of the process.
elapsedapp	float	The elapsed user time of the process, in seconds.
elapsedsys	float	The elapsed system time of the process, in seconds.
elapsedreal	float	The elapsed wall-clock time of the process, in seconds.
rcode	integer	The return code of the process.
retrycount	integer	The value of the task retry counter when the invocation ran. A task retry or job restart will increment it if the task had become active at some point.
resumecount	integer	A number ordering an iteration by its resume pass. iteration number of times ( the resume of the task retry counter when the invocation ran. A task retry or job restart will increment it if the task had become active at some point.
resumable	boolean	A boolean value that is true if the command can be resumed.
bladeid	id	The bladeid of the blade the invocation's command is running or ran on.

## Blade Attributes

name	type	description
name	string	The blade name. It is a unique identifier of the blade, and is typically the blade's hostname.
ipaddr	string	The IP address of the host.
port	integer	The number of the port on which the blade is listening.
osname	string	The name of the operating system on which the blade is running.

osversion	string	The version of operating system on which the blade is running.
boottime	timestamp	The boot time of the host.
numcpu	integer	The number of cpus/cores of the host.
loadavg	float	The CPU load average of the host.
availmemory	float	The available memory of the host, in GB.
availdisk	float	The available disk space of the host, in GB.
version	string	The Tractor blade version.
profile	string	The profile name used by the blade.
nimby	string	The NIMBY status of the blade.
starttime	timestamp	The start time of the blade process.
numslots	integer	The total number of slots on the blade.
udi	float	The Universal Desirability Index of the blade, which helps certain blades be assigned tasks sooner than other blades. A higher value means the higher the chance of the blade being assigned work.
status	string	A status note for the blade.
heartbeattime	timestamp	The time the blade last contacted the engine.
bladeid	id	A universally unique identifier for a blade. Over time, attributes such as hostname, port, and IP address may be reused by different hosts. This id is used when populating invocation records so that the particular blade that ran the command can be properly identified in perpetuity.
cleartime	timestamp	The time a user requested the blade data be cleared. Only blades with a heartbeattime greater than the cleartime will be visible in the dashboard.
gpulabel	string	Blade host GPU information.

## Overriding Attributes

Jobs and tasks can specify attribute values that pertain to commands. These attributes are *service*, *tags*, *maxslots*, and *envkey*. Values defined at the command level override those specified at the task and job level; values defined at the task level override those specified at the job level.