


Engine Metrics


The "Engine Metrics" tab provides a graphical recent history of several engine health and performance statistics. Tractor engine periodically samples these values and sends them to subscribed Dashboard clients. New samples are taken approximately every 15 seconds, with approximately 10 minutes of history displayed in the graphs, by default.

These graphs are intended to provide administrators with a quick overview of the engine's status, and to help identify performance issues as they arise. Since the graphs are based on samples taken from a fast-changing system, not all of the displayed values will exactly match the system state on a moment by moment basis.




waiting tasks: the number the tasks waiting to be executed across the entire job queue.

active tasks: the number of tasks that have been dispatched to blades and are currently running.




total slots: the sum of "slot" counts reported by all blades on the farm. Slots are an abstract unit of assignable capacity. Some sites will configure one slot per host, some will use one slot per core, others will add their own plug-ins to the blade to generate dynamic slot values.

active slots: the number of slots consumed by actively running tasks. The number used by each command is specified in the job script, and may be a dynamic number between the given "at-least" and "at-most" values.



requests / sec: an estimate of the total number of HTTP requests received by tractor-engine per second. This includes all requests from blades, browsers, and external scripts.

dispatches / sec: the rate of successful task assignments. This is the number of blade requests for work that are successfully matched with waiting tasks from the job queue, and delivered back to the blade for execution, per second.




intake backlog: an internal measure of the inbound HTTP requests to tractor-engine that have been received but not yet delegated to a specific engine subsystem for processing.


assigner backlog: the number of pending requests from blades waiting to be assigned tasks from the queue. Some of these requests may be successfully processed but not result in task launch, such as when there are no remaining queued tasks or when the blade's service keys do not match the requirements of any waiting commands.

i/o pool backlog: tractor uses a pool of "shipper" threads to concurrently handle i/o operations on connected sockets. This value is a count of the high-level i/o actions that are waiting for an available thread.

monitor backlog: a count of the UI-specific HTTP requests that are waiting for processing.



engine load: the normalized current "load average" on the tractor-engine host itself. The displayed value is the system one-minute loadavg, such as given by 'uptime' or 'top', but then normalized by the number of CPU cores on the host. Thus a 16-core host with 8 fully utilized cores would report a nominal loadavg of "8.0", whereas this graph will report "0.5". The intent is to indicate the overall fraction of the host's CPU capacity that is in use. This load may not always be due to entirely tractor-engine activity. Other activity on the host, such as file-serving or other running applications will also contribute to its load.



user connections: the number of recently connected client sessions. This includes Dashboard sessions as well as other clients that are also using the HTTP request protocol to receive job status updates or to request job or blade details. Since the "end" of an ongoing HTTP "session" can be difficult to characterize, the engine considers a connection to still be "current" until a relatively long interval expires without additional activity from the expected client. Specifically, fully reloading a tractor browser window temporarily cause both the old and new sessions to be counted as potentially active.