

## Arvados - Feature #14021

### [crunch-dispatch-slurm] option to set job priority directly instead of using nice values

08/13/2018 06:14 PM - Tom Clegg

<b>Status:</b>	New	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>			
<b>Description</b>			
<p>Background: crunch-dispatch-slurm uses slurm's "nice" feature to make its slurm jobs' priority order match the priority order of their corresponding arvados containers. This is better than setting slurm job priorities directly in that it doesn't require c-d-s to have slurm administrator privileges. However, on older versions of slurm, its effectiveness is limited by slurm's maximum nice value of 10000.</p> <p>crunch-dispatch-slurm could offer a configuration flag that causes it to set slurm job priority directly ("scontrol update jobid=X priority=Y") instead of restricting itself to non-privileged "nice" operations.</p> <p>This configuration would be useful on systems running slurm 15 where it is acceptable to give slurm administrator privileges to crunch-dispatch-slurm.</p> <p>Possible implementation strategies:</p> <ol style="list-style-type: none"><li>1. Assume Arvados jobs are the only ones running. Choose a high priority like 2^32-3 (slurm's usual "TOP" priority) for the highest-priority Arvados job, and count down from there using the existing PrioritySpread logic.</li><li>2. Wait for each submitted job to appear in the slurm queue, and note its automatically assigned priority. When updating priorities, start at the highest auto-assigned priority of all jobs still shown in queue, and count down from there using the existing PrioritySpread logic.<ul style="list-style-type: none"><li>◦ Caveat: on a cloud installation, jobs frequently drop into "admin hold" state (priority=0), so it might be some time before c-d-s gets a chance to see a given job's automatically assigned priority. The resulting behavior might be confusing.</li></ul></li></ol>			

#### History

##### #2 - 08/13/2018 06:15 PM - Tom Clegg

- Target version set to *To Be Groomed*

##### #3 - 08/14/2018 02:00 PM - Tom Clegg

- Subject changed from *[crunch-dispatch-slurm] Config option to set job priority directly instead of using nice values* to *[crunch-dispatch-slurm] option to set job priority directly instead of using nice values*

##### #4 - 07/06/2021 09:09 PM - Peter Amstutz

- Target version deleted (*To Be Groomed*)