

## Arvados - Bug #13991

### crunch-dispatch-slurm does not warn when slurm MaxJobCount reached

08/08/2018 11:08 AM - Joshua Randall

<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>	Crunch	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	To Be Groomed		
<b>Description</b>			
SLURM has a default MaxJobCount of 10000. MaxJobCount is not specified in the example slurm.conf suggested by Arvados install docs ( <a href="https://doc.arvados.org/install/crunch2-slurm/install-slurm.html">https://doc.arvados.org/install/crunch2-slurm/install-slurm.html</a> ). Perhaps it should be, so it is clear that this parameter might matter to Arvados.			
When hitting the MaxJobCount, I would have expected crunch-dispatch-slurm to log some sort of warning indicating that it was unable to queue all of the jobs because of the limit. I did not see it saying anything that seemed to mean that.			

### History

#### #1 - 08/09/2018 03:44 PM - Tom Morris

- Target version set to To Be Groomed

#### #2 - 08/09/2018 07:46 PM - Tom Clegg

Tested on 9tee4 with MaxJobCount=2.

```
$ sbatch -N1 <(printf '#!/bin/sh\nsleep 8000\n')
sbatch: error: Slurm temporarily unable to accept job, sleeping and retrying.
^C
```

It would be more convenient to get a non-zero exit code, so we don't have to scrape stderr while sbatch is still running.

There's an sbatch --immediate option that fails if the allocation can't be granted, but that's not what we want either. We want to fail only if the job can't be *queued*.

So it seems the solution is for crunch-dispatch-slurm to monitor stderr while sbatch is running, and if that message appears:

- Log a suggestion to increase MaxJobCount
- Avoid starting more sbatch processes until this one exits (incidentally, we only recently stopped serializing *all* sbatch invocations).

We should also log any other unexpected messages from sbatch, to make other similar problems easier to diagnose.