

## Arvados - Story #12289

### Migrate to Slurm API

09/19/2017 08:31 PM - Tom Morris

<b>Status:</b> New	<b>Start date:</b> 09/26/2017
<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> Arvados Future Sprints	
<b>Description</b> Currently nodemanager & crunch-dispatch-slurm parse the output of Slurm command line utilities to get the information that it needs. This is fragile and labor intensive since changes to the output format must be tracked by the parsing code.  We should investigate switching to using the actual Slurm API. <a href="https://github.com/PySlurm/pyslurm">https://github.com/PySlurm/pyslurm</a>	
<b>Subtasks:</b> Task # 12336: Investigate Slurm API <span style="float: right;"><b>New</b></span>	

#### History

#1 - 09/26/2017 06:56 PM - Tom Morris

- Description updated

#2 - 09/26/2017 07:55 PM - Tom Clegg

Potential benefits include better compatibility across slurm versions (less CLI-output-parsing code) and lower overhead (less shelling out to CLI).

However, <https://github.com/PySlurm/pyslurm> says

You will need to instruct the setup.py script where either the Slurm install root directory or where the Slurm libraries and Slurm include files are

It sounds like this package provides an API client for a specific version of slurm, which could be just another distribution/deployment nightmare. It only makes sense if every (current+future) slurm library version is forward & backward compatible with all slurm versions. Otherwise, CLI tools are more reliable -- at least it's *possible* to accommodate versions with "if v3 then X else Y" when necessary.