

## Arvados - Feature #4935

### [Crunch] Choose worker nodes where the desired docker image is already available.

01/08/2015 04:48 PM - Bryan Cosca

<b>Status:</b>	New	<b>Start date:</b>	01/08/2015
<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>	Arvados Future Sprints		

#### Description

The scenario where user A is running jobs and debugging a pipeline to work on arvados. User A ran his job and it ran on compute 1. Then he realized something failed, and five minutes later he fixes the problem and commits the change. User A then runs his job again and it runs again on compute1 and fails again and fixes the problem. Then, he runs his job again and realizes it doesn't fail right away. The compute node has changed from under him and he is now using compute2. He now has to wait ~7 minutes for his docker image to load into the node or he has to wait ~10 minutes for the node to spin up and then load his image. His ~1 minute turn around time for development has just become ~10 minutes, leaving the user to simply wait.

If its possible, I would like to develop on one node and not have these nodes pulled from under me if I'm not using them for ~5-10 minutes.

#### History

##### #1 - 01/13/2015 08:43 PM - Tom Clegg

- Subject changed from Dedicated Nodes to users to [Crunch] Choose worker nodes where the desired docker image is already available.

##### #2 - 01/19/2015 02:52 PM - Brett Smith

- Category set to Crunch

##### #3 - 01/19/2015 07:32 PM - Bryan Cosca

A really good example imo is this instance: [qr1hi-d1hrv-sbjgb8za1dpp5kc](#). The first two jobs ran fine on compute16 and then for some reason, the third job chose compute43 which had #5006 problem.

##### #4 - 01/20/2015 07:30 PM - Tom Clegg

- Target version set to Arvados Future Sprints