

## Arvados - Feature #15457

### [Controller] Delegate new container requests to other clusters based on location of input data

07/10/2019 04:59 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>	API	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Arvados Future Sprints		
<b>Description</b>			
<p>When a client creates a new container request (and doesn't specify a desired cluster ID) controller should resolve all input collection to PDHs as needed, and then:</p> <ul style="list-style-type: none"><li>• if all inputs are available locally, create a local container request (as the current implementation does in all cases)</li><li>• otherwise, rank local/remote clusters according to how many of the input data bytes they have on hand, and execute a "create CR" request on the highest-ranking cluster -- being sure to specify the chosen cluster ID so the remote cluster doesn't have to repeat the ranking/selection process itself.<ul style="list-style-type: none"><li>◦ if the local cluster is tied with a remote, choose the local cluster</li><li>◦ use the <code>file_size_total</code> collection attribute</li></ul></li></ul> <p>At least for now, don't go to too much trouble to be precise -- if a mount only refers to a small file in a large collection, it's OK to rank by the entire collection size.</p> <p>If a remote cluster returns an error during the "probe for inputs" phase, drop that cluster from the list of candidates.</p> <p>If a remote cluster returns an error or times out when submitting a container request, fall back to submitting to the local cluster (unless this fallback is disabled via config knob). If this fallback is enabled, the remote call should time out in 1/2 the remaining portion of the locally configured API request timeout (see <a href="#">Deadline()</a>). If the local cluster fails (whether or not a remote has also been attempted), just return the error to the caller.</p> <p>Add an entry to the CR's properties hash indicating how the cluster was chosen, including any errors encountered when probing or submitting to remotes.</p>			
<b>Related issues:</b>			
Related to Arvados - Bug #14710: [Workbench] Child containers run on federate...			<b>New</b>

#### History

##### #1 - 07/10/2019 05:14 PM - Tom Clegg

- Related to Bug #14710: [Workbench] Child containers run on federated clusters do not show up added

##### #2 - 07/10/2019 05:29 PM - Tom Clegg

- Description updated

##### #3 - 07/10/2019 05:33 PM - Tom Clegg

- Description updated

##### #4 - 07/10/2019 05:34 PM - Tom Clegg

- Description updated

##### #5 - 07/10/2019 05:44 PM - Tom Clegg

- Description updated

##### #6 - 07/10/2019 05:44 PM - Tom Morris

- Target version changed from To Be Groomed to Arvados Future Sprints

- Story points set to 3.0