

Arvados - Story #14996

Log crunch-run resource usage statistics

03/19/2019 05:56 PM - Tom Morris

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Arvados Future Sprints		
Description			
Currently we have crunchstat.txt for container stats and hoststat.txt for node level stats, but we have no way of recording the time and resources for file upload by crunch-run after the container exits. This would be very useful since, for I/O intensive jobs like bcl2fastq, this can represent as much time as the compute phase of the job.			

History

#1 - 03/19/2019 06:17 PM - Peter Amstutz

In order to collect fine-grained metrics, one option is for crunch-run to have a bootstrap mode that creates a cgroup and runs crunch-run and arv-mount inside it.

#2 - 03/20/2019 04:40 PM - Eric Biagiotti

- Story points set to 2.0

Go cgroup module should make this easier <https://godoc.org/github.com/containerd/cgroups>. Might be worth trying to have crunch-run move itself to a cgroup in bootstrap mode. If that doesn't work, we'd have to start another process.

Once crunch-run and arv-mount are in their own cgroup, the arvados crunchstat go module can automatically interpret stats.

#3 - 04/03/2019 02:07 PM - Tom Morris

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