

Arvados - Bug #4882

[Crunch] crunchstat reports surprising CPU usage when container appears and disappears

12/29/2014 08:51 PM - Bryan Cosca

Status:	Resolved	Start date:	12/29/2014
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	Crunch	Estimated time:	0.00 hour
Target version:	2015-05-20 sprint		
Description			
Background			
<p>source:services/crunchstat/crunchstat.go has a list of locations where cgroup accounting files are likely to be found. Historically this was needed because the location varies from one kernel/system to another. Each time crunchstat reads stats, it tries each location in turn until it finds one that exists. When the target container is up, this works fine, but the last two locations give accounting information for the host instead of the container. Often, crunchstat tries to read accounting data before the container has come up or after the container has been destroyed (in the interval before it gets notified that the docker process has exited). Naturally, the transition between host stats and container stats results in reporting wild deltas.</p> <p>The transition from one stat file to another causes an "error" log message which is easy to misinterpret as an error affecting the job, or explaining some later job failure -- see #5523.</p>			
Proposed fix			
One or more of:			
<ul style="list-style-type: none">• Don't report host-level stats if container-level stats have ever been reported.• Don't report host-level stats ever.• When the previous sample point and the current sample point come from different accounting files, don't print a delta.			
Bug report			
Originally reported as "Log graph goes out of bounds":			
log_graph_oob.png about 5/6th of the way through, you can see the yellow point go off the screen. Suggestion: maybe the graph should scale in real time... Hmm, I'm not sure because then the graph would keep scaling and would probably get annoying. It would still be useful to at least see that point though.			
Subtasks:			
Task # 5932: Review 4882-no-host-when-container			Resolved
Related issues:			
Related to Arvados - Bug #5523: [Crunch] crunchstat should not report errors ...		New	05/07/2015

Associated revisions

Revision 31d8ece6 - 05/08/2015 06:42 AM - Tom Clegg

Merge branch '4882-no-host-when-container' closes #4882

History

#1 - 01/07/2015 09:25 PM - Tim Pierce

- Target version set to Bug Triage

#2 - 01/16/2015 08:07 PM - Tim Pierce

- Subject changed from Log graph goes out of bounds to [Workbench] Log graph goes out of bounds

- Category set to Workbench

#3 - 01/19/2015 07:40 PM - Brett Smith

The job in question is [qr1hi-8i9sb-ajgkcn4cuk4cntm](#). At 2014-12-29_21:48:34, where previous CPU user values were 10 or below, suddenly it

skyrocketed to over 4000. After that, samples were below 1 for the remainder of the job.

crunchstat apparently didn't miss a sample. The logs show them being reported every 10 seconds consistently.

I wouldn't rule out the possibility that this is a bug in the underlying data—either in crunchstat, or the records it reads.

#4 - 01/21/2015 08:39 PM - Tom Clegg

- Subject changed from [Workbench] Log graph goes out of bounds to [Workbench] crunchstat reports surprising CPU usage when container appears and disappears

- Description updated

- Category changed from Workbench to Crunch

- Target version changed from Bug Triage to Arvados Future Sprints

#5 - 01/21/2015 08:39 PM - Tom Clegg

- Subject changed from [Workbench] crunchstat reports surprising CPU usage when container appears and disappears to [Crunch] crunchstat reports surprising CPU usage when container appears and disappears

#6 - 05/05/2015 11:09 PM - Tom Clegg

- Description updated

#7 - 05/06/2015 02:00 PM - Tom Clegg

- Description updated

#8 - 05/07/2015 02:55 AM - Tom Clegg

- Target version changed from Arvados Future Sprints to 2015-05-20 sprint

#9 - 05/07/2015 02:55 AM - Tom Clegg

- Status changed from New to In Progress

#10 - 05/07/2015 05:20 PM - Peter Amstutz

It seems likely that there will be spurious "stats not available" messages during container startup and teardown, it would be nice to suppress those since that is an expected condition. What do you think about adding a counter or timer so that it only prints the notice after several consecutive failed attempts to read the stats?

#11 - 05/07/2015 05:44 PM - Peter Amstutz

- Assigned To set to Tom Clegg

#12 - 05/07/2015 05:58 PM - Tom Clegg

Peter Amstutz wrote:

It seems likely that there will be spurious "stats not available" messages during container startup and teardown, it would be nice to suppress those since that is an expected condition. What do you think about adding a counter or timer so that it only prints the notice after several consecutive failed attempts to read the stats?

That seems reasonable. "Several" should probably be two: neither of those races lasts more than a few milliseconds unless something's wrong. It would be ideal to handle the startup race by waiting a short time and repolling, to avoid missing the first interval of data. Added this to [#5523](#) ("verbose logs sound too scary") rather than blocking this ("stats are incorrect").

#13 - 05/08/2015 06:45 AM - Tom Clegg

- Status changed from In Progress to Resolved

Applied in changeset arvados|commit:31d8ece649c63b008ed79a930a5b237ec795ff22.

Files

log_graph_oob.png	56.3 KB	12/29/2014	Bryan Cosca
-------------------	---------	------------	-------------