

Arvados - Bug #12690

crunchstat-summary --container-request does not report on keep i/o

11/28/2017 03:17 PM - Bryan Cosca

Status:	Resolved	Start date:	12/05/2017
Priority:	Normal	Due date:	
Assigned To:	Peter Amstutz	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	2018-10-17 sprint		
Description			
Crunch-run reports the arv-mount stats in arv-mount.txt in the logging collection, but crunchstat-summary doesn't know to look there.			

Associated revisions

Revision 35e38be6 - 04/30/2018 08:37 PM - Tom Clegg

Fix wrong event_type when getting buffered live logs for crunch2.

refs #12690

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

Revision 04b9d6e9 - 10/05/2018 01:36 AM - Peter Amstutz

Merge branch '12690-12748-crunchstat-summary'

refs #12690

refs #12748

refs #13759

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <pamstutz@veritasgenetics.com>

History

#1 - 11/28/2017 03:17 PM - Bryan Cosca

- Subject changed from *crunchstat-summary does not report on keep i/o* to *crunchstat-summary --container-request does not report on keep i/o*

#2 - 11/28/2017 03:49 PM - Peter Amstutz

- Description updated

#3 - 11/28/2017 06:07 PM - Tom Morris

- Target version set to *To Be Groomed*

#4 - 11/29/2017 07:20 PM - Tom Morris

- Target version changed from *To Be Groomed* to *2017-12-20 Sprint*

#5 - 12/04/2017 03:35 PM - Tom Morris

- Assigned To set to *Tom Morris*

#6 - 12/05/2017 12:08 AM - Tom Morris

- Status changed from *New* to *In Progress*

#7 - 12/05/2017 03:06 PM - Peter Amstutz

- Target version changed from *2017-12-20 Sprint* to *2017-12-06 Sprint*

#8 - 12/05/2017 04:27 PM - Peter Amstutz

The arv-mount.txt stats file was just added recently.

Unfortunately this means if you run crunchstat-summary on an older run, you get lots of errors:

```
Exception in thread Thread-402:
```

```
Traceback (most recent call last):
  File "/usr/lib/python2.7/threading.py", line 801, in __bootstrap_inner
    self.run()
  File "/usr/lib/python2.7/threading.py", line 754, in run
    self.__target(*self.__args, **self.__kwargs)
  File "/home/peter/work/scripts/venv/local/lib/python2.7/site-packages/crunchstat_summary-0.1.20171205000254-py2.7.egg/crunchstat_summary/summarizer.py", line 526, in run_and_release
    return target(*args, **kwargs)
  File "/home/peter/work/scripts/venv/local/lib/python2.7/site-packages/crunchstat_summary-0.1.20171205000254-py2.7.egg/crunchstat_summary/summarizer.py", line 74, in run
    self._run(logdata)
  File "/home/peter/work/scripts/venv/local/lib/python2.7/site-packages/crunchstat_summary-0.1.20171205000254-py2.7.egg/crunchstat_summary/summarizer.py", line 78, in _run
    for line in logdata:
  File "/home/peter/work/scripts/venv/local/lib/python2.7/site-packages/crunchstat_summary-0.1.20171205000254-py2.7.egg/crunchstat_summary/reader.py", line 33, in __iter__
    self._readers.append(collection.open(filename))
  File "/home/peter/work/scripts/venv/local/lib/python2.7/site-packages/arvados_python_client-0.1.20171010180436-py2.7.egg/arvados/collection.py", line 660, in open
    raise IOError(errno.ENOENT, "File not found", path)
IOError: [Errno 2] File not found: 'arv-mount.txt'
```

It needs to handle missing stats files.

#9 - 12/06/2017 07:06 AM - Tom Morris

Is the job/collection buried some place in that stack trace? I'm not seeing it.

#10 - 12/06/2017 07:58 PM - Tom Morris

- Target version changed from 2017-12-06 Sprint to 2017-12-20 Sprint

#11 - 01/02/2018 06:07 PM - Tom Morris

- Target version changed from 2017-12-20 Sprint to 2018-01-17 Sprint

#12 - 01/17/2018 07:37 PM - Tom Morris

- Target version changed from 2018-01-17 Sprint to 2018-01-31 Sprint

#13 - 02/01/2018 08:54 PM - Tom Morris

- Target version changed from 2018-01-31 Sprint to 2018-02-14 Sprint

#14 - 02/14/2018 08:19 PM - Tom Morris

- Target version changed from 2018-02-14 Sprint to 2018-02-28 Sprint

#15 - 02/28/2018 08:14 PM - Tom Morris

- Target version changed from 2018-02-28 Sprint to 2018-03-14 Sprint

#16 - 03/14/2018 07:08 PM - Tom Morris

- Target version changed from 2018-03-14 Sprint to 2018-03-28 Sprint

#17 - 03/28/2018 03:15 PM - Tom Morris

- Target version changed from 2018-03-28 Sprint to 2018-04-11 Sprint

#18 - 04/11/2018 03:22 PM - Tom Morris

- Target version changed from 2018-04-11 Sprint to 2018-04-25 Sprint

#19 - 04/25/2018 02:50 PM - Tom Morris

- Target version changed from 2018-04-25 Sprint to 2018-05-09 Sprint

#20 - 05/09/2018 03:20 PM - Tom Morris

- Target version changed from 2018-05-09 Sprint to 2018-05-23 Sprint

#21 - 05/23/2018 02:45 PM - Tom Morris

- Target version changed from 2018-05-23 Sprint to 2018-06-06 Sprint

#22 - 06/06/2018 02:44 PM - Tom Morris

- Target version changed from 2018-06-06 Sprint to 2018-06-20 Sprint

#23 - 06/20/2018 03:21 PM - Tom Morris

- Target version changed from 2018-06-20 Sprint to 2018-07-03 Sprint

#24 - 07/03/2018 03:04 PM - Tom Morris

- Target version changed from 2018-07-03 Sprint to 2018-07-18 Sprint

#25 - 07/18/2018 03:06 PM - Tom Morris

- Target version changed from 2018-07-18 Sprint to 2018-08-01 Sprint

#26 - 08/01/2018 03:14 PM - Tom Morris

- Target version changed from 2018-08-01 Sprint to 2018-08-15 Sprint

#27 - 08/15/2018 03:17 PM - Tom Morris

- Target version changed from 2018-08-15 Sprint to 2018-09-05 Sprint

#28 - 08/17/2018 02:45 PM - Peter Amstutz

It's not looking for the "crunchstat:" prefix on the arv-mount log lines and so it is trying to parse error messages:

```
Error parsing value u'of' (stat u'memory', category u'out'): ValueError("invalid literal for int() with base 10: 'of'",)
2018-08-16T03:21:35.090831222Z out of memory
```

```
Error parsing value u'of' (stat u'memory', category u'out'): ValueError("invalid literal for int() with base 10: 'of'",)
2018-08-16T03:22:04.094908800Z out of memory
```

#29 - 09/05/2018 02:51 PM - Tom Morris

- Target version changed from 2018-09-05 Sprint to 2018-09-19 Sprint

#30 - 09/05/2018 08:05 PM - Tom Morris

Finally getting back to looking at this. arv-mount.txt doesn't use the crunchstat: prefix. Was this arv-mount.txt or crunchstat.txt? Do you have a link to the job so that I can use it for testing? (I can probably work around it even without the prefix).

#31 - 09/05/2018 08:11 PM - Peter Amstutz

Tom Morris wrote:

Finally getting back to looking at this. arv-mount.txt doesn't use the crunchstat: prefix. Was this arv-mount.txt or crunchstat.txt? Do you have a link to the job so that I can use it for testing? (I can probably work around it even without the prefix).

Every line in this file has "crunchstat:" after the timestamp and before the actual stats.

https://collections.e51c5.arvadosapi.com/c=10e2b6c05210cf58317b8d904c02f729-489/_arv-mount.txt?disposition=inline

#32 - 09/19/2018 03:19 PM - Tom Morris

- Target version changed from 2018-09-19 Sprint to 2018-10-03 Sprint

#33 - 10/03/2018 03:20 PM - Peter Amstutz

- Target version changed from 2018-10-03 Sprint to 2018-10-17 sprint

#34 - 10/03/2018 03:20 PM - Peter Amstutz

- Assigned To changed from Tom Morris to Peter Amstutz

#35 - 10/05/2018 01:39 AM - Peter Amstutz

- Status changed from In Progress to Resolved

#36 - 11/13/2018 08:53 PM - Tom Morris

- Release set to 14