

Arvados - Bug #15734

[a-d-c] needs to populate node.json in the container log collection

10/17/2019 08:23 PM - Ward Vandewege

Status:	Resolved	Start date:	10/22/2019
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2019-11-20 Sprint		

Description

Node manager drops a file called node.json in the container log collection, with an output like this:

```
{
  "created_at": "2017-04-12T04:37:11.016719000Z",
  "crunch_worker_state": "busy",
  "domain": "qrlhi.arvadosapi.com",
  "etag": "xyrne3ax1barycf21f0xw49i",
  "first_ping_at": "2019-06-01T18:53:52.714980000Z",
  "hostname": "compute3",
  "href": "/nodes/qrlhi-7ekkf-z6lokkesrtdw9ia",
  "ip_address": "10.26.64.21",
  "job_uuid": null,
  "kind": "arvados#node",
  "last_ping_at": "2019-06-01T18:56:02.166756000Z",
  "modified_at": "2019-06-01T18:56:02.170027000Z",
  "modified_by_client_uuid": null,
  "modified_by_user_uuid": "qrlhi-tpzed-0000000000000000",
  "nameservers": [
    "10.26.0.11"
  ],
  "owner_uuid": "qrlhi-tpzed-0000000000000000",
  "properties": {
    "cloud_node": {
      "price": 0.192,
      "size": "Standard_D4s_v3"
    },
    "total_cpu_cores": 4,
    "total_ram_mb": 16018,
    "total_scratch_mb": 32747
  },
  "slot_number": 3,
  "status": "running",
  "uuid": "qrlhi-7ekkf-z6lokkesrtdw9ia"
}
```

This information originates from the Arvados node record. That table is no longer used by a-d-c. We need the 'properties' section to be available in that file for container cost accounting. A-d-c should populate the node.json file like this:

```
{
  "properties": {
    "cloud_node": {
      "price": 0.192,
      "size": "Standard_D4s_v3"
    },
    "total_cpu_cores": 4,
    "total_ram_mb": 16018,
    "total_scratch_mb": 32747
  },
}
```

instead of just dropping an empty file in the container logs. If it makes sense to change the node.json format that is fine, as long as the new format is easily machine readable.

Subtasks:		
Task # 15741: Review 15734-dispatchcloud-node-info		Resolved
Related issues:		
Related to Arvados - Story #15759: [arvados-dispatch-cloud] deploy/run correc...	Resolved	12/30/2019

Associated revisions

Revision 4cad483 - 10/24/2019 06:22 PM - Tom Clegg

Merge branch '15734-dispatchcloud-node-info'

fixes #15734

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

Revision e87f5931 - 10/29/2019 06:22 PM - Tom Clegg

Merge branch '15734-dispatchcloud-node-info'

fixes #15734

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

History

#1 - 10/17/2019 08:27 PM - Ward Vandewege

- Description updated

#2 - 10/17/2019 08:29 PM - Ward Vandewege

- Target version changed from To Be Groomed to 2019-11-06 Sprint

#3 - 10/17/2019 08:31 PM - Ward Vandewege

- Description updated

#4 - 10/21/2019 07:41 PM - Tom Clegg

- Status changed from New to In Progress

- Assigned To set to Tom Clegg

15734-dispatchcloud-node-info @ [9481ff4a22314c0d5acffe78fbc7595278414e6f](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1606/) --
<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1606/>

crunch-run copies the relevant InstanceTypes config entry into node.json if it was dispatched via arvados-dispatch-cloud. This means the format differs depending on whether you're using crunch-dispatch-slurm or arvados-dispatch-cloud.

#5 - 10/24/2019 04:21 PM - Peter Amstutz

Tom Clegg wrote:

15734-dispatchcloud-node-info @ [9481ff4a22314c0d5acffe78fbc7595278414e6f](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1606/) --
<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1606/>

crunch-run copies the relevant InstanceTypes config entry into node.json if it was dispatched via arvados-dispatch-cloud. This means the format differs depending on whether you're using crunch-dispatch-slurm or arvados-dispatch-cloud.

```
+ // In order to stay compatible with recent dev/experimental
+ // versions of crunch-run, we need to pass a map with string
+ // values only, so we JSON-encode the instance type
+ // record. Once worker images are updated we can skip this and
+ // just pass {"InstanceType": wkr.instType}.
```

I found this comment quite confusing. I think what you saying is that versions of crunch-run *prior to this one* will fail if given a non-string value, so the workaround is for dispatcher to pass the json text nested in a string and write it directly to the file instead of decoding and reencoding it.

This reminds me that the idea of provisioning a known copy of crunch-run on newly booted nodes would neatly avoid these kinds of compatibility problems. We discussed this, is there actually a ticket for that?

Rest LGTM.

#6 - 10/24/2019 06:15 PM - Tom Clegg

- Related to Story #15759: [arvados-dispatch-cloud] deploy/run correct version of crunch-run binary on worker nodes added

#7 - 10/24/2019 07:31 PM - Tom Clegg

- Status changed from In Progress to Resolved

- % Done changed from 0 to 100

Applied in changeset [arvados|4cad4832ed4dd58e841628232c472e99b6f7d1b](#).

#8 - 10/24/2019 08:49 PM - Tom Clegg

Note in order for this to work, you need to update crunch-run on the worker image, as well as updating arvados-dispatch-cloud.

#9 - 10/29/2019 03:01 PM - Tom Clegg

- Status changed from Resolved to In Progress

Reopening. Current implementation doesn't work because

- a-d-c passes InstanceType map on stdin of crunch-run
- crunch-run reads stdin into a map[string]interface{} and saves string values to environment
- crunch-run starts a child crunch-run -no-detach; child inherits string values from the map via environment vars -- but not the map[string]interface{} itself
- crunch-run (parent) exits
- crunch-run (child) looks for InstanceType in runner.dispatchEnv, which is always empty because it's only populated in the parent process

#10 - 10/29/2019 03:35 PM - Tom Clegg

15734-dispatchcloud-node-info @ [f5f4c5f6088707c86c43ca2bd70a70c14699e450](#) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1615/>

- abandon hope of passing as a struct in future
- when logging, use the env var (which is propagated to the detached crunch-run child) instead of a map (which isn't)

#11 - 10/29/2019 06:19 PM - Peter Amstutz

Tom Clegg wrote:

15734-dispatchcloud-node-info @ [f5f4c5f6088707c86c43ca2bd70a70c14699e450](#) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1615/>

- abandon hope of passing as a struct in future
- when logging, use the env var (which is propagated to the detached crunch-run child) instead of a map (which isn't)

Darn, the interaction with "detach" was not obvious to either of us.

I wish we had a better way to test this than rolling it out to staging clusters (eg [#15370](#))

Technically, couldn't we pipe dispatchEnv to the detached process the same way the parent gets it from c-d-c?

However, I think the simpler approach of propagating the plain string is probably better in this case.

LGTM (but I have no way of realistically testing it)

#12 - 10/29/2019 07:20 PM - Tom Clegg

- Status changed from In Progress to Resolved

Applied in changeset [arvados|e87f5931f41937cf210cdebd42783c9dc8c8c723](#).

#13 - 11/01/2019 05:58 PM - Ward Vandewege

Tom Clegg wrote:

Applied in changeset [arvados|e87f5931f41937cf210cdebd42783c9dc8c8c723](#).

Confirmed working, see https://workbench.qr1hi.arvadosapi.com/container_requests/qr1hi-xvhdp-w4xk395ewnppx07#Log

#14 - 11/11/2019 09:50 PM - Tom Morris

- Status changed from Resolved to In Progress

- Target version changed from 2019-11-06 Sprint to 2019-11-20 Sprint

The output in the collection that Ward pointed to isn't compatible with the previous format and won't work with any of the existing cost scripts. The key naming, structure, and units have all changed.

New:

```
{
  Name: "Standard_D1_v2",
  ProviderType: "Standard_D1_v2",
  VCPUs: 1,
  RAM: 3584000000,
  Scratch: 50000000000,
  IncludedScratch: 50000000000,
  AddedScratch: 0,
  Price: 0.057,
  Preemptible: false
}
```

Original:

```
{
  "properties": {
    "cloud_node": {
      "price": 0.192,
      "size": "Standard_D4s_v3"
    },
    "total_cpu_cores": 4,
    "total_ram_mb": 16018,
    "total_scratch_mb": 32747
  },
}
```

#15 - 11/19/2019 09:36 PM - Ward Vandewege

- Status changed from In Progress to Resolved

Tom Morris wrote:

The output in the collection that Ward pointed to isn't compatible with the previous format and won't work with any of the existing cost scripts. The key naming, structure, and units have all changed.

Yes, that's by design. The old format was simply a dump of our (idiosyncratic) arvados node object.

We don't use those objects anymore with a-d-c. The format was undocumented anyway.

The new format is much more straightforward and includes more information, which was not available before. We only had internal tools consuming it anyway, so this is a good cleanup.

Oh, and TomC did check with the reviewer (me) about this before changing the format. Resolving this ticket again.

#16 - 01/22/2020 02:09 PM - Peter Amstutz

- Release set to 22