

Arvados - Bug #12055

[node manager] ec2 set tags on create

08/01/2017 01:42 PM - Peter Amstutz

Status: Resolved	Start date: 08/16/2017
Priority: Normal	Due date:
Assigned To: Lucas Di Pentima	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2017-09-13 Sprint	
Description https://issues.apache.org/jira/browse/LIBCLOUD-930 ec2 node_create(ex_metadata) uses TagSpecification.N → https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_RunInstances.html Fix this in our local fork (https://github.com/curoverse/libcloud) https://dev.arvados.org/issues/11953#note-10 Fix node manager to use node_create(ex_metadata) instead of ex_create_tags(). Update tests. Contribute PR to libcloud upstream (https://github.com/apache/libcloud) Update libcloud tests.	
Subtasks:	
Task # 12066: Review 12055-nodemanager-ec2-tags	Resolved
Related issues:	
Related to Arvados - Story #8999: [node manager] Upgrade to libcloud 2.0	Closed
Related to Arvados - Bug #9223: [Node manager] Uses huge amount of RAM on AWS	Resolved

Associated revisions

Revision 2ac65228 - 09/01/2017 10:04 PM - Lucas Di Pentima

12055: Merge branch '12055-nodemanager-ec2-tags'
Closes #12055

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <ldipentima@veritasgenetics.com>

History

#1 - 08/01/2017 01:42 PM - Peter Amstutz

- Subject changed from [node manager] ec2 set tags on create to [node manager] ec2 set tags on create

#2 - 08/01/2017 01:42 PM - Peter Amstutz

- Description updated

#3 - 08/02/2017 05:40 AM - Tom Morris

- Target version set to 2017-08-16 sprint

I'm pretty sure that this is a duplicate and that someone (Lucas?) was already working on fixing it, but I'll add it to the upcoming sprint while I investigate.

#4 - 08/02/2017 07:25 PM - Peter Amstutz

- Description updated

#5 - 08/02/2017 07:25 PM - Peter Amstutz

- Description updated

#6 - 08/02/2017 07:28 PM - Tom Morris

- Assigned To set to Lucas Di Pentima

#7 - 08/02/2017 07:28 PM - Peter Amstutz

- Description updated

- Assigned To deleted (Lucas Di Pentima)

#8 - 08/02/2017 07:28 PM - Peter Amstutz

- Assigned To set to Lucas Di Pentima

#9 - 08/02/2017 07:32 PM - Peter Amstutz

- Description updated

#10 - 08/02/2017 07:33 PM - Peter Amstutz

- Description updated

#11 - 08/10/2017 02:11 PM - Lucas Di Pentima

- Status changed from New to In Progress

#12 - 08/16/2017 04:01 AM - Lucas Di Pentima

Updates at [f5f91a293](#)

Test run: <https://ci.curoverse.com/job/developer-run-tests/416/>

- Updated our libcloud fork (<https://github.com/curoverse/libcloud/tree/apache-libcloud-0.20.2.dev4>) so that tags are passed to the create node call.
- Updated nodemanager to provide ex_metadata to the create_node() call, instead of assigning them after the creation operation.
- Updated tests and dependencies.

#13 - 08/16/2017 03:27 PM - Peter Amstutz

I'm not sure this is correct:

```
params['TagSpecification.1.ResourceType'] = 'instance'
params['TagSpecification.1.Tags'] = [
    {'Key': k, 'Value': v} for k, v in tags.items()
]
```

The documentation (https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_RunInstances.html) seems to suggest it should be:

```
params['TagSpecification.0'] = {
    'ResourceType': 'instance',
    'Tags': [{'Key': k, 'Value': v} for k, v in tags.items()]
}
```

#14 - 08/16/2017 03:42 PM - Peter Amstutz

Aha, here's an example raw request:

```
https://ec2.amazonaws.com/?Action=RunInstances
&ImageId=ami-31814f58
&InstanceType=t2.large
&MaxCount=2
&MinCount=1
&KeyName=my-key-pair
&SubnetId=subnet-b2a249da
&TagSpecification.1.ResourceType=instance
&TagSpecification.1.Tag.1.Key=webserver
&TagSpecification.1.Tag.1.Value=production
&TagSpecification.2.ResourceType=volume
&TagSpecification.2.Tag.1.Key=cost-center
&TagSpecification.2.Tag.1.Value=cc123
&AUTHPARAMS
```

#15 - 08/16/2017 03:53 PM - Peter Amstutz

I think you can do this in a one liner:

```
if not 'ex_metadata' in create_kwargs:
    create_kwargs['ex_metadata'] = {}
```

→

```
create_kwargs.setdefault('ex_metadata', {})
```

#16 - 08/16/2017 04:27 PM - Lucas Di Pentima

Updated libcloud code at commit 79ec53df - branch apache-libcloud-0.20.2.dev4

Updated nodemanager at [f2019e704](#)

Test run: <https://ci.curoverse.com/job/developer-run-tests/417/>

- Libcloud: fixed the way parameters are built as per the provided examples
- Nodemanager
 - Merged master so the sdk/cwl tests don't fail
 - Changed the code following Peter's suggestion (oneliner comment)

#17 - 08/16/2017 06:38 PM - Lucas Di Pentima

- Target version changed from 2017-08-16 sprint to 2017-08-30 Sprint

#18 - 08/17/2017 01:31 PM - Peter Amstutz

apache-libcloud-0.20.2.dev4 @ commit:79ec53df225c98f9e9c5ce67e32715c91b06daeb LGTM

12055-nodemanager-ec2-tags @ [f2019e7042d12088bce45f8c2ad52ec600a4076d](#) LGTM

We need to test this on 4xphq.

#19 - 08/17/2017 01:36 PM - Peter Amstutz

Suggested test strategy:

1. log in to 4xphq
2. sv stop arvados-node-manager
3. virtualenv nodemanager-venv
4. . nodemanager-venv/bin/activate
5. git clone libcloud && python setup.py install
6. git clone arvados && python setup.py install node manager
7. arvados-node-manager /etc/arvados-node-manager/config.ini
8. Submit a job & watch the logs at /etc/sv/arvados-node-manager/log/main/current
9. Ask Nico and Javier to use the ec2 console to confirm that the tags are being set.

#20 - 08/18/2017 02:30 PM - Lucas Di Pentima

I have been working on this since yesterday afternoon.

Things accomplished:

- Accessed as root to 4xphq. Check
- Created a virtualenv on /root/lucas-tests/arvados-node-manager/
- Installed a-n-m & libcloud dev4 version inside that virtualenv
- Stopped system a-n-m and tried to run the test one using --foreground so I could see its logging

Things still missing:

- When running the test nodemanager, I started to get error messages from Amazon like this one:

```
2017-08-17 18:40:49 ComputeNodeSetupActor.32d525e9b2ca[16149] INFO: Sending create_node request for node size Compute Optimized Large Instance.
2017-08-17 18:40:49 ComputeNodeSetupActor.32d525e9b2ca[16149] WARNING: Re-raising error (no retry): UnknownParameter: The parameter TagSpecification is not recognized
Traceback (most recent call last):
  File "/root/lucas-tests/arvados-node-manager/local/lib/python2.7/site-packages/arvados_node_manager-0.1.20170816162215-py2.7.egg/arvnodeman/computenode/__init__.py", line 81, in retry_wrapper
    ret = orig_func(self, *args, **kwargs)
  File "/root/lucas-tests/arvados-node-manager/local/lib/python2.7/site-packages/arvados_node_manager-0.1.20170816162215-py2.7.egg/arvnodeman/computenode/dispatch/__init__.py", line 133, in create_cloud_node
    self.arvados_node)
  File "/root/lucas-tests/arvados-node-manager/local/lib/python2.7/site-packages/arvados_node_manager-0.1.20170816162215-py2.7.egg/arvnodeman/computenode/driver/__init__.py", line 184, in create_node
    raise create_error
BaseHTTPError: UnknownParameter: The parameter TagSpecification is not recognized
```

- In the middle of my testing, when trying to do changes on libcloud and/or nodemanager and restart it, 4xphq's api server started to give this

errors to nodemanager:

```
2017-08-18_14:03:33.61977 2017-08-18 14:03:33 googleapiclient.http[19766] WARNING: Invalid JSON content from response: {"errors":["Forbidden"],"error_token":"1503065013+08989b6e"}
2017-08-18_14:03:33.62011 2017-08-18 14:03:33 ComputeNodeSetupActor.7c41e7e722f2[19766] WARNING: Client error: <HttpError 403 when requesting https://4xphq.arvadosapi.com/arvados/v1/nodes?alt=json returned "Forbidden"> - scheduling retry in 180 seconds
2017-08-18_14:03:33.62011 Traceback (most recent call last):
2017-08-18_14:03:33.62012   File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/__init__.py", line 81, in retry_wrapper
2017-08-18_14:03:33.62012       ret = orig_func(self, *args, **kwargs)
2017-08-18_14:03:33.62013   File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/dispatch/__init__.py", line 116, in create_arvados_node
2017-08-18_14:03:33.62013       self.arvados_node = self._arvados.nodes().create(body={}).execute()
2017-08-18_14:03:33.62013   File "/usr/lib/python2.7/dist-packages/oauth2client/util.py", line 140, in positional_wrapper
2017-08-18_14:03:33.62013       return wrapped(*args, **kwargs)
2017-08-18_14:03:33.62013   File "/usr/lib/python2.7/dist-packages/googleapiclient/http.py", line 840, in execute
2017-08-18_14:03:33.62013       raise HttpError(resp, content, uri=self.uri)
2017-08-18_14:03:33.62014 ApiError: <HttpError 403 when requesting https://4xphq.arvadosapi.com/arvados/v1/nodes?alt=json returned "Forbidden">
```

- I even reestablished the system's nodemanager but this kind of messages still occur.
- Now on 4xphq's workbench, it seems to be 1 busy node and 1 pipeline instance pending.

#21 - 08/18/2017 09:17 PM - Lucas Di Pentima

I have been trying without success to get from AWS docs which EC2 API version is the one that starts including the TagSpecification parameter. Our libcloud fork uses this one: '2013-10-15'

I tried changing it to the one that libcloud is currently using: '2016-11-15', and the result was that no more "Unknown parameter" error happened, but I tried to run the test suite and there are a lot of tests failing, so I'm not confident that this would not cause issues.

Right now with this nodemanager+libcloud combo, 4xphq seems to be doing work but there's something that don't seem normal: there are 2 idle nodes and 1 busy, are those 2 idle nodes supposed to be shutdown?

I'm going on vacations for a week in a couple of hours, I'll shutdown the nodemanager just in case, remember that the stable one for some reason has an invalid token and since yesterday wasn't able to communicate with the api server.

The test version is on /root/lucas-tests/

#22 - 08/30/2017 11:17 AM - Lucas Di Pentima

Created a new branch on our libcloud fork, based against upstream version 2.2: <https://github.com/curoverse/libcloud/tree/apache-libcloud-2.2.0.dev1>
Updated nodemanager dependency at [83e428528](https://github.com/curoverse/libcloud/pull/142)

As we talked yesterday, I've tested libcloud 2.2 locally with our nodemanager test suite, and got no errors, so I made a 2.2 branch on our fork and applied the TagSpecification patch.

Before we deploy to 4xphq for further testing, the complete test suite is running at: <https://ci.curoverse.com/job/developer-run-tests/419/>

#23 - 08/30/2017 03:34 PM - Lucas Di Pentima

Stress test run successful at: https://workbench.4xphq.arvadosapi.com/pipeline_instances/4xphq-d1hrv-mspsxk8tgxfmwhd

#24 - 08/30/2017 06:46 PM - Lucas Di Pentima

- Target version changed from 2017-08-30 Sprint to 2017-09-13 Sprint

#25 - 08/31/2017 03:14 PM - Lucas Di Pentima

Running the entire test suite against the latest libcloud 2.2.0.dev1 changes: <https://ci.curoverse.com/job/developer-run-tests/422/>

#26 - 08/31/2017 07:26 PM - Peter Amstutz

If the stable release of libcloud is 2.2.0, we need our fork to be 2.2.1.dev1. Right now we have a package filename "2.2.0.dev1" but package metadata has version "2.2.0" and Python doesn't like that.

Note that "dev" versions are considered *before* non-dev versions (in other words 2.2.0.dev1 < 2.2.0 < 2.2.1.dev1 < 2.2.1).

Please create a 2.2.1.dev1 github branch and update the package version & node manager & build script version references accordingly.

This also avoids the pull request problem, since the PR is originating from 2.2.0.dev1 (the libcloud PR can't mess with the upstream version number).

#27 - 08/31/2017 07:45 PM - Peter Amstutz

Also, make sure to test package building (this is how I noticed the versioning problems above).

```
$ export WORKSPACE=$HOME/arvados
$ cd arvados/build
$ ./run-build-packages-one-target.sh --target debian8 --only-build python-apache-libcloud
$ ./run-build-packages-one-target.sh --target debian8 --only-build arvados-node-manager
$ ./run-build-packages-one-target.sh --target debian8 --only-test arvados-node-manager
```

#28 - 08/31/2017 10:11 PM - Lucas Di Pentima

Updates at [b13876638cf57b400bd59513a0a1811b3d2993a1](https://github.com/Arvados/arvados/pull/1811)

Created new 2.2.1.dev1 branch including the latest fixes that were requested on the PR and updated dependencies on build scripts. Tested package creation successfully, but the last package testing returns this error:

```
START: arvados-node-manager test on arvados/package-test:debian8
Reading package lists...
Building dependency tree...
Reading state information...
Some packages could not be installed. This may mean that you have
requested an impossible situation or if you are using the unstable
distribution that some required packages have not yet been created
or been moved out of Incoming.
The following information may help to resolve the situation:

The following packages have unmet dependencies:
 arvados-node-manager : Depends: python-arvados-python-client (>= 0.1.20170731145219) but it is not going to be installed
                       Depends: python-future but it is not installable
E: Unable to correct problems, you have held broken packages.
ERROR: arvados-node-manager test on arvados/package-test:debian8 failed with exit status 100
Failed package tests: arvados-node-manager
```

#29 - 09/01/2017 02:30 PM - Lucas Di Pentima

Update: Building all the packages made the package test succeed

#30 - 09/01/2017 03:15 PM - Nico César

review @ [b13876638cf57b400bd59513a0a1811b3d2993a1](https://github.com/Arvados/arvados/pull/1811)

we have to be able to package the libcloud and check that we're able to install in all distributions without dependency problems (including jessie for example), include the proper changes to the build/build.list to make FPM create the package for us

#31 - 09/01/2017 04:07 PM - Lucas Di Pentima

The script build/run-build-packages-one-target.sh already makes the deb package from our libcloud fork. I tried to install it on my debian8 dev instance:

```
lucas@curoverse:~/arvados_local$ sudo dpkg -i packages/debian8/python-apache-libcloud_2.2.1.dev1-2_all.deb
Selecting previously unselected package python-apache-libcloud.
(Reading database ... 104618 files and directories currently installed.)
Preparing to unpack .../python-apache-libcloud_2.2.1.dev1-2_all.deb ...
Unpacking python-apache-libcloud (2.2.1.dev1-2) ...
Setting up python-apache-libcloud (2.2.1.dev1-2) ...
lucas@curoverse:~/arvados_local$
```

#32 - 09/01/2017 07:56 PM - Lucas Di Pentima

Testing libcloud-2.2.1.dev1 & a-n-m from [b13876638cf57b400bd59513a0a1811b3d2993a1](https://github.com/Arvados/arvados/pull/1811) on 4xphq showed a high RAM usage that made nodemanager to die at bootup of memory starvation.

This may be related to [#9223](#) & [#12163](#):

```
2017-09-01_18:31:23.26201 2017-09-01 18:31:23 root[542] INFO: /usr/bin/arvados-node-manager 0.1.20170831201516
, libcloud 2.2.1.dev1
2017-09-01_18:31:23.30414 2017-09-01 18:31:23 requests.packages.urllib3.connectionpool[542] DEBUG: Starting ne
w HTTPS connection (1): ec2.us-east-1.amazonaws.com
2017-09-01_18:31:33.48756 2017-09-01 18:31:33 requests.packages.urllib3.connectionpool[542] DEBUG: https://ec2
.us-east-1.amazonaws.com:443 "GET /?SignatureVersion=2&AWSAccessKeyId=AKIAJCNUIVXKTYNJ5OSQ&Timestamp=2017-09-0
1T18%3A31%3A23Z&SignatureMethod=HmacSHA256&Version=2016-11-15&Signature=dEawnGnFX1Wkr3SEnPHhei5FZAFKJ2dSc77KYw
h7A%2FU%3D&Action=DescribeImages HTTP/1.1" 200 None
2017-09-01_18:32:09.74370 2017-09-01 18:32:09 root[542] ERROR: Uncaught exception during setup
2017-09-01_18:32:09.74371 Traceback (most recent call last):
2017-09-01_18:32:09.74372   File "/usr/lib/python2.7/dist-packages/arvnodeman/launcher.py", line 131, in main
2017-09-01_18:32:09.74372     server_calculator = build_server_calculator(config)
```

```

2017-09-01_18:32:09.74372 File "/usr/lib/python2.7/dist-packages/arvnodeman/launcher.py", line 74, in build_
server_calculator
2017-09-01_18:32:09.74372     cloud_size_list = config.node_sizes(config.new_cloud_client().list_sizes())
2017-09-01_18:32:09.74373 File "/usr/lib/python2.7/dist-packages/arvnodeman/config.py", line 129, in new_clo
ud_client
2017-09-01_18:32:09.74373     driver_class=driver_class)
2017-09-01_18:32:09.74373 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/ec2.py", line
60, in __init__
2017-09-01_18:32:09.74373     driver_class)
2017-09-01_18:32:09.74373 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/__init__.py",
line 71, in __init__
2017-09-01_18:32:09.74373     new_pair = init_method(self.create_kwargs.pop(key))
2017-09-01_18:32:09.74374 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/ec2.py", line
63, in _init_image_id
2017-09-01_18:32:09.74374     return 'image', self.search_for(image_id, 'list_images')
2017-09-01_18:32:09.74374 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/__init__.py",
line 120, in search_for
2017-09-01_18:32:09.74375     term, list_method, key, **kwargs)
2017-09-01_18:32:09.74375 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/__init__.py",
line 102, in search_for_now
2017-09-01_18:32:09.74376     items = list_func(**kwargs)
2017-09-01_18:32:09.74376 File "/usr/lib/python2.7/dist-packages/libcloud/compute/drivers/ec2.py", line 3535
, in list_images
2017-09-01_18:32:09.74376     self.connection.request(self.path, params=params).object
2017-09-01_18:32:09.74376 File "/usr/lib/python2.7/dist-packages/arvnodeman/computenode/driver/ec2.py", line
30, in request
2017-09-01_18:32:09.74376     return super(ANMEC2Connection, self).request(*args, **kwargs)
2017-09-01_18:32:09.74376 File "/usr/lib/python2.7/dist-packages/libcloud/common/base.py", line 637, in requ
est
2017-09-01_18:32:09.74377     response = responseCls(**kwargs)
2017-09-01_18:32:09.74377 File "/usr/lib/python2.7/dist-packages/libcloud/common/base.py", line 159, in __in
it__
2017-09-01_18:32:09.74377     self.object = self.parse_body()
2017-09-01_18:32:09.74377 MemoryError

```

I've monkey-patched nodemanager passing an owner filter 'self' to libcloud's list_images() call so it doesn't retrieve all the available images, and now it starts without issues.

Running a stress test to see if it works ok.

#33 - 09/01/2017 08:12 PM - Lucas Di Pentima

Updates at [1ba39510d](#)

Test run: <https://ci.curoverse.com/job/developer-run-tests/425/>

Retrieve the node image list from AWS but filtering those that are owned by us, to avoid high memory usage. (See [#9223](#), [#12163](#))

Stress test with --disable-reuse seems to be going ok: https://workbench.4xphq.arvadosapi.com/pipeline_instances/4xphq-d1hrv-909j9zpf1vjxsou

#34 - 09/01/2017 10:10 PM - Anonymous

- Status changed from In Progress to Resolved

- % Done changed from 0 to 100

Applied in changeset arvados|commit:2ac65228ff6b32921e6c8194b6c51ce9a710f385.