

Arvados - Feature #14716

[keep-web] Use cluster config file

01/10/2019 06:50 PM - Peter Amstutz

Status: Resolved	Start date: 08/08/2019
Priority: Normal	Due date:
Assigned To: Lucas Di Pentima	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2019-08-14 Sprint	
Description	
Subtasks:	
Task # 15391: Review 14716-webdav-cluster-config	Resolved
Task # 15543: Review 14716-doc-fixes	Resolved
Related issues:	
Related to Arvados - Story #13648: [Epic] Use one cluster configuration file ...	Resolved

Associated revisions

Revision 1ea75125 - 08/06/2019 07:20 PM - Tom Clegg

Merge branch '14716-anonymous-token'

refs #14716

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

Revision 0fb04af5 - 08/13/2019 02:32 PM - Lucas Di Pentima

Merge branch '14716-webdav-cluster-config'

Refs #14716

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

Revision b6499675 - 08/13/2019 06:21 PM - Lucas Di Pentima

Merge branch '14716-doc-fixes'

Closes #14716

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

History

#1 - 01/10/2019 06:50 PM - Peter Amstutz

- Related to Story #13648: [Epic] Use one cluster configuration file for all components added

#2 - 02/06/2019 07:23 PM - Tom Morris

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

- Story points set to 1.0

#3 - 06/19/2019 03:31 PM - Tom Morris

- Target version changed from Arvados Future Sprints to 2019-07-03 Sprint

#4 - 06/19/2019 03:35 PM - Lucas Di Pentima

- Assigned To set to Lucas Di Pentima

#5 - 06/27/2019 07:22 PM - Lucas Di Pentima

- Status changed from New to In Progress

#6 - 07/03/2019 02:43 PM - Lucas Di Pentima

- Target version changed from 2019-07-03 Sprint to 2019-07-17 Sprint

#7 - 07/17/2019 02:33 PM - Lucas Di Pentima

- Target version changed from 2019-07-17 Sprint to 2019-07-31 Sprint

#8 - 07/31/2019 02:52 PM - Lucas Di Pentima

- Target version changed from 2019-07-31 Sprint to 2019-08-14 Sprint

#9 - 08/06/2019 01:16 PM - Lucas Di Pentima

Updates at [2c7b49e87](https://ci.curoverse.com/job/developer-run-tests/1448/) - branch 14716-webdav-cluster-config
Test run: <https://ci.curoverse.com/job/developer-run-tests/1448/>

There's a keep-web integration test that's failing and I'm not being able to tell why. It's the one testing access to a collection using /collections/uuid/path. Supposedly the anonymous token is used on that case ignoring all other provided credentials, I made sure that the anonymous user token is configured, however but it still fails. Some Workbench integration tests also fail, and they're related to the download feature so maybe I'm seeing the same issue from 2 different points of view.

#10 - 08/06/2019 06:35 PM - Tom Clegg

It looks like the test was relying on a bug. The "/collections/\$uuid/path" form is only supposed to use the configured anonymous tokens, and ignore caller-supplied tokens. But when no anonymous tokens were configured (as in this test), keep-web was incorrectly using the supplied credentials.

14716-anonymous-token @ [f89544af7f3d38bd61b4216527d66897eb08dcd0](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1451/) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1451/>

- don't use supplied credentials at /collections/\$uuid/path, even if no anon tokens are configured
- fix test → expect 404

#11 - 08/08/2019 03:18 PM - Lucas Di Pentima

Updates at [a33badcdc](https://ci.curoverse.com/job/developer-run-tests/1459/) - branch 14716-webdav-cluster-config
Test run: <https://ci.curoverse.com/job/developer-run-tests/1459/>

- Merged Tom's fix so now keep-web tests pass OK
- Updated arvbox to make keep-web work
- Migrated deprecated NodeProfiles on arvbox cluster config

There're some workbench integration tests still failing. I'm not being able to run them correctly on my local dev VM (even on master, they fail in other ways) so not sure how to debug this. I think what's happening is a configuration issue between keep-web and workbench of some kind. Already tried the new keep-web on arvbox and seem to work fine.

#12 - 08/08/2019 05:02 PM - Lucas Di Pentima

The workbench integration tests that are failing are for example (from <https://ci.curoverse.com/job/developer-run-tests-apps-workbench-integration/1533/console>):

```
13:55:05 56) Failure:
13:55:05 DownloadTest#test_download_anonymous_content_from_keep-web_by_portable_data_hash [/tmp/workspace/developer-run-tests-apps-workbench-integration/apps/workbench/test/integration/download_test.rb:84]:
13:55:05 --- expected
13:55:05 +++ actual
13:55:05 @@ -1,2 +1 @@
13:55:05 -"Hello world
13:55:05 -"
13:55:05 +nil
13:55:05
13:55:05 57) Failure:
13:55:05 DownloadTest#test_download_anonymous_content_from_keep-web_by_uuid [/tmp/workspace/developer-run-tests-apps-workbench-integration/apps/workbench/test/integration/download_test.rb:84]:
13:55:05 --- expected
13:55:05 +++ actual
13:55:05 @@ -1,2 +1 @@
13:55:05 -"Hello world
13:55:05 -"
13:55:05 +nil
13:55:05
13:55:05 58) Failure:
13:55:05 DownloadTest#test_download_from_keep-web_using_a_session_token [/tmp/workspace/developer-run-tests-apps-workbench-integration/apps/workbench/test/integration/download_test.rb:84]:
13:55:05 Expected: "w a z"
13:55:05 Actual: nil
```

Checking the keep-web.log from the artifacts (

<https://ci.curoverse.com/job/developer-run-tests-apps-workbench-integration/1533/artifact/tmp/keep-web.log>) I'm seeing that all requests that mention the 'Hello World.txt' and 'w a z' files were served correctly, that's why I'm thinking there has to be a problem in the interaction between keep-web and workbench of some sort.

#13 - 08/08/2019 06:52 PM - Tom Clegg

config docs

Set `Services.Controller.Insecure: true` if your API server's TLS certificate is not signed by a recognized CA.

should be `TLS.Insecure: true`

```
"http://:9002/": {}
```

should be `"http://keep_web_hostname_goes_here:9002/": {}`

```
exec sudo -u nobody keep-web -config=/path/to/arvados.yml
```

I wonder if we need a note that `arvados.yml` needs to be world-readable?

I wonder if we should drop `-config=/path/to/arvados.yml` from the install docs and just rely on the default, since we already instructed the reader to use the default location.

to your Workbench cluster configuration file

remove Workbench

```
ExternalURL: "https://download.<span class='userinput'>uuid_prefix</span>.your.domain/c=%{uuid_or_pdh}"
```

This shouldn't have the `%{uuid_or_pdh}` stuff any more -- Workbench now adds the `/c=` stuff (or puts the UUID in the host part) depending on whether the URL has a wildcard.

```
ExternalURL: "https://%{uuid_or_pdh}--collections.<span class='userinput'>uuid_prefix</span>.your.domain"
```

Likewise these should be `"https://*--collections.uuid_prefix.your.domain"` etc.

code

`AnonymousUserToken:`

Looks like you added a second copy of this key to `config.default.yml`?

```
(ldr *Loader) LoadDefaults()
```

I think I already rejected this in [#14717#note-10](#). You can do this instead:

```
cfg, err := config.NewLoader(bytes.NewBufferString("Clusters: {zzzz: {} }"), nil).Load()
```

If you want a helper function for that, it should be in `arvadostest` -- cluster `zzzz` is only for testing.

```
yaml "gopkg.in/yaml.v2"
```

should use `ghodss/yaml` like everywhere else

```
DefaultConfig()
```

this seems to have changed purpose -- rename to `newConfig(*arvados.Config)` and update comment?

```
TestLegacyConfig
```

This is in `services/keep-web`, but seems to be testing code in `lib/config` -- maybe it should move to `lib/config`?

The rest LGTM.

#22 - 08/13/2019 06:40 PM - Lucas Di Pentima

- *Status changed from In Progress to Resolved*

Applied in changeset [arvados|b64996755741a71c4f02cf3ef9ea2b7ba1d0e2e1](#).

#23 - 01/21/2020 09:20 PM - Peter Amstutz

- *Release set to 22*