

## Arvados - Feature #15531

### [SDK] Migrate federation to central LoginCluster

08/07/2019 02:27 PM - Peter Amstutz

<b>Status:</b>	Resolved	<b>Start date:</b>	09/23/2019
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>	Peter Amstutz	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2019-10-23 Sprint		
<b>Description</b>			
Mass user migration to centralized federated user uuids.			
<ol style="list-style-type: none"><li>1. Generate a list of all existing users across clusters (fill in the "home" column with "LoginCluster" for every user)</li><li>2. Produce report listing each email address, existing user uuid, if the user owns anything, the cluster that the user will be migrated to</li><li>3. Admin reviews report</li><li>4. Report is fed back in</li><li>5. If no local user exists for a home cluster, create a new user with the email address</li><li>6. Use user merge to redirect all the old user accounts (and reassign all their data) to user on the home cluster. Decide where to relocate ownership (user or a project owned by user) based on which accounts are empty or not.</li></ol>			
Note: API server needs to allow migrating objects owned by remote users to federated user, and make the old remote user disappear (currently migrating remote accounts is not supported) (this was specifically requested by the customer).			
API merge method allows specifying the owner_uuid			
Write some tests (based on mocks) for the tool.			
<b>Subtasks:</b>			
Task # 15617: Review 15531-logincluster-migrate			<b>Resolved</b>
<b>Related issues:</b>			
Related to Arvados - Story #15529: [API] [Controller] Share user account data...		<b>Resolved</b>	<b>08/22/2019</b>
Related to Arvados - Task #15208: Migration script		<b>Resolved</b>	<b>04/18/2019</b>
Related to Arvados - Feature #15061: Redirect users to log in with correct fe...		<b>Resolved</b>	<b>04/18/2019</b>
Related to Arvados - Story #15577: [API] User ownership reassign		<b>Resolved</b>	<b>11/08/2019</b>
Related to Arvados - Story #15720: [API] Unified user listing across all clus...		<b>Resolved</b>	<b>11/19/2019</b>

#### Associated revisions

##### Revision 0c92112a - 10/08/2019 07:05 PM - Peter Amstutz

arvbox update config.yml when config override is updated refs #15531

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <[pamstutz@veritasgenetics.com](mailto:pamstutz@veritasgenetics.com)>

##### Revision 7ab77ac9 - 10/11/2019 03:04 PM - Peter Amstutz

Merge branch '15531-logincluster-migrate' refs #15531

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <[pamstutz@veritasgenetics.com](mailto:pamstutz@veritasgenetics.com)>

#### History

##### #1 - 08/07/2019 02:27 PM - Peter Amstutz

- Status changed from New to In Progress

##### #2 - 08/07/2019 02:27 PM - Peter Amstutz

- Subject changed from User migration tool to User migration tool for federation

- Description updated

- Status changed from In Progress to New

**#3 - 08/07/2019 02:28 PM - Peter Amstutz**

- Related to Story #15529: [API] [Controller] Share user account database with a group of trusted clusters added

**#4 - 08/07/2019 02:29 PM - Peter Amstutz**

- Description updated

**#5 - 08/07/2019 03:36 PM - Peter Amstutz**

- Description updated

**#6 - 08/07/2019 03:39 PM - Peter Amstutz**

- Subject changed from User migration tool for federation to User management/migration tool

**#7 - 08/07/2019 03:44 PM - Peter Amstutz**

- Description updated

**#8 - 08/21/2019 03:33 PM - Peter Amstutz**

- Description updated

**#9 - 08/21/2019 05:09 PM - Tom Morris**

- Related to Task #15208: Migration script added

**#10 - 08/21/2019 05:09 PM - Tom Morris**

- Related to Feature #15061: Redirect users to log in with correct federated identity added

**#11 - 08/21/2019 05:22 PM - Peter Amstutz**

- Description updated

**#12 - 08/21/2019 05:23 PM - Peter Amstutz**

- Description updated

**#13 - 08/21/2019 05:24 PM - Tom Morris**

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

- Story points set to 3.0

**#14 - 08/22/2019 06:23 PM - Peter Amstutz**

- Subject changed from User management/migration tool to [SDK] Migrate federation to central LoginCluster

**#15 - 09/10/2019 02:33 PM - Peter Amstutz**

- Target version changed from Arvados Future Sprints to 2019-09-11 Sprint

**#16 - 09/10/2019 02:34 PM - Peter Amstutz**

- Assigned To set to Peter Amstutz

**#17 - 09/11/2019 01:36 PM - Peter Amstutz**

Make sure to update documentation

**#18 - 09/11/2019 03:01 PM - Peter Amstutz**

- Target version changed from 2019-09-11 Sprint to 2019-09-25 Sprint

**#19 - 09/18/2019 02:55 PM - Tom Clegg**

- Related to Story #15577: [API] User ownership reassign added

**#20 - 09/18/2019 02:56 PM - Tom Clegg**

When merging user A into user B without redirect (e.g., Admin is taking over Bob's account when Bob leaves), make sure permission links pointing to B don't end up pointing to A.

#21 - 09/23/2019 09:18 PM - Peter Amstutz

15531-logincluster-migrate @ [858008ed900a92bf7dcf2e7b14b3162b3d17ab03](#)

<https://ci.curoverse.com/job/developer-run-tests/1544/>

- API: enable merging remote users into local ones
- API: merging failed to rename git repositories with new username, fixed
- API: setting user redirect clears 'username' so it is possible for the target user to take it over
- API: add "exported config" endpoint to discovery document
- API: Fix remote token validate to use RemoteHosts.\*.Insecure
- fed-migrate script: don't require tokens.csv if the cluster has LoginCluster set
- fed-migrate script: report conflicting usernames
- fed-migrate script: migration adjusts username
- fed-migrate script: migrate all instances of user (local and remote) across all clusters
- fed-migrate script: always fetch user list, can re-run the report and won't do anything (and won't break) the 2nd time

See `sdk/python/tests/fed-migrate/README` for information about integration tests.

#22 - 09/24/2019 01:41 PM - Peter Amstutz

- Status changed from New to In Progress

#23 - 09/24/2019 06:43 PM - Lucas Di Pentima

Some preliminary comments and questions:

- Is it necessary to commit .cwl files that are generated from .cwllex?
- Typo at `services/api/app/models/user.rb:L304`
- Running the integration tests:
  - The 3 `arvbox build+run` took ~80 minutes on a VM with 8 GB dedicated RAM + some swap. If this is meant to be run on Jenkins maybe we can speed up the build step by building a new instance and then copy it over to the other 2 to avoid downloading/recompiling the same things 3 times.
  - The test run failed with the following message:

```
[...]
curl --fail --insecure --silent https://172.17.0.4:8000/discovery/v1/apis/arvados/v1/rest
+ export ARVADOS_API_HOST=172.17.0.4:8000
+ arvbox cat /var/lib/arvados/superuser_token
+ export ARVADOS_API_TOKEN=1fj54h1c9allcf8f0q2m9v61l2oht8ebgbpwnu0magylk6jw09
+ export ARVADOS_API_HOST_INSECURE=1
+ arvbox cat /var/lib/arvados/vm-uuid
+ ARVADOS_VIRTUAL_MACHINE_UUID=xg6ao-2x53u-lu55ry9hbni9p7p
+ python -c import arvados ; arvados.api().virtual_machines().get( uuid='xg6ao-2x53u-lu55ry9hbni9p7p' ).execute()
+ sleep 3
+ python -c import arvados ; arvados.api().virtual_machines().get( uuid='xg6ao-2x53u-lu55ry9hbni9p7p' ).execute()
+ sleep 3
[...]
+ python -c import arvados ; arvados.api().virtual_machines().get( uuid='xg6ao-2x53u-lu55ry9hbni9p7p' ).execute()
+ sleep 3
+ python -c import arvados ; arvados.api().virtual_machines().get( uuid='xg6ao-2x53u-lu55ry9hbni9p7p' ).execute()
INFO [job main_2_embed_2_3] Max memory used: 59MiB
INFO [job main_2_embed_2_3] completed success
INFO [step main_2_embed_2_3] completed success
INFO [workflow main_2_3] starting step superuser_tok_3_3
INFO [step superuser_tok_3_3] start
INFO [job superuser_tok_3_3] /tmp/bg9nxacx$ arvbox \
  cat \
  /var/lib/arvados/superuser_token > /tmp/bg9nxacx/superuser_token.txt
INFO [job superuser_tok_3_3] completed success
INFO [step superuser_tok_3_3] completed success
INFO [workflow main_2_3] completed success
INFO [step main_2] completed success
INFO [workflow ] starting step run_test_3
INFO [step run_test_3] start
INFO [workflow run_test_3] start
INFO [workflow run_test_3] starting step main_1_2
INFO [step main_1_2] start
INFO [job main_1] /tmp/viuxpr0b$ python \
  /tmp/tmpy0wk_8ny/stg963ab711-df8b-4e69-8c46-63c272e77fa2/create_users.py \
  _script
Traceback (most recent call last):
```

```

File "/tmp/tmpy0wk_8ny/stg963ab711-df8b-4e69-8c46-63c272e77fa2/create_users.py", line 23, in <module>
    arvados.api(host=j["arvados_api_hosts"][1], token=make_token(newtok), insecure=True).users().current().execute()
File "/home/lucas/venv-arvados/lib/python3.5/site-packages/googleapiclient/_helpers.py", line 130, in positional_wrapper
    return wrapped(*args, **kwargs)
File "/home/lucas/venv-arvados/lib/python3.5/site-packages/googleapiclient/http.py", line 840, in execute
    raise HttpError(resp, content, uri=self.uri)
arvados.errors.ApiError: <HttpError 401 when requesting https://172.17.0.3:8000/arvados/v1/users/current?alt=json returned "Not logged in (req-iliiv1fgh6smtfu7vmsxd)">
INFO [job main_1] Max memory used: 38MiB
WARNING [job main_1] completed permanentFail
WARNING [step main_1_2] completed permanentFail
INFO [workflow run_test_3] completed permanentFail
WARNING [step run_test_3] completed permanentFail
INFO [workflow ] completed permanentFail
{
  "supertok": [
    "3wipyhlee0hdzuet42co9h0s21e41vmfzhsqs1vleid9r5ngca",
    "3q6c9potn0l7khs6ehe45bj5kssm16d8c366hlyfsqjsupow6w",
    "1fj54h1c9allcf8f0q2m9v61l2oht8ebgpbwnu0magylk6jw09"
  ],
  "report": null
}
WARNING Final process status is permanentFail

```

#### #24 - 09/24/2019 10:18 PM - Lucas Di Pentima

Some more comments:

- "hotreset" documentation is missing on arvbox's help output
- I've tried to manually test the script by doing:
  - Removed LoginCluster entries on fedbox2 & fedbox3 (and did hotreset on them)
  - Created the same admin account on the 3 clusters
  - Appended their certs on my VM's system ca-certificates.crt file
  - Extracted the tokens to a CSV file and ran the script with --tokens and --report, got the following error message:

```

$ arv-federation-migrate --tokens ~/tokens.csv --report -
Reading /home/lucas/tokens.csv
Contacting 172.17.0.2:8000
Contacting 172.17.0.3:8000
Contacting 172.17.0.4:8000
Checking that the federation is well connected
Traceback (most recent call last):
  File "/home/lucas/venv-arvados/bin/arv-federation-migrate", line 4, in <module>
    __import__('pkg_resources').run_script('arvados-python-client==1.4.1.20190923192457', 'arv-federation-migrate')
  File "/home/lucas/venv-arvados/lib/python3.5/site-packages/pkg_resources/__init__.py", line 739, in run_script
    self.require(requires)[0].run_script(script_name, ns)
  File "/home/lucas/venv-arvados/lib/python3.5/site-packages/pkg_resources/__init__.py", line 1494, in run_script
    exec(code, namespace, namespace)
  File "/home/lucas/venv-arvados/lib/python3.5/site-packages/arvados_python_client-1.4.1.20190923192457-py3.5.egg/EGG-INFO/scripts/arv-federation-migrate", line 7, in <module>
    main()
  File "/home/lucas/venv-arvados/lib/python3.5/site-packages/arvados_python_client-1.4.1.20190923192457-py3.5.egg/arvados/commands/federation_migrate.py", line 54, in main
    config = arv.configs().get().execute()
AttributeError: 'Resource' object has no attribute 'configs'

```

- Just in case, I removed the discovery doc cache files on ~/.cache/arvados/ but got the same result.

#### #25 - 09/24/2019 10:20 PM - Lucas Di Pentima

Hmmm, now that I think about it, those fedboxN instances may be based on master, tomorrow I'll make sure they're based on the correct branch.

#### #26 - 09/25/2019 01:05 PM - Lucas Di Pentima

Update: I've restarted the fedboxN instances with the correct branch and the federation migrate tests exited with success.

#### #27 - 09/25/2019 03:18 PM - Peter Amstutz

- Target version changed from 2019-09-25 Sprint to 2019-10-09 Sprint

## #28 - 09/25/2019 08:23 PM - Lucas Di Pentima

More comments & questions while I continue reviewing:

- File `sdk/python/arvados/commands/federation_migrate.py`
  - At line 34 I got an `IndexError` exception when using a tokens file with an empty line (the last line), maybe it would be nice to ignore those.
  - There's commented out code at line 60
  - Seems that there's repeated code blocks at lines 112-121 & 125-135.
    - I'm not sure what they do, maybe add some comments?
    - If the first for-loop is to get `homeuid` to be empty string or some unique uid on a list, wouldn't be something like the following be more clear?:

```
uuids = [x["uuid"] for x in accum]
homeuid = (len(set(uuids)) == 1) and uuids[0] or ""
```

- Do you think it is convenient to add a `--version` argument (and maybe report the version to stdout on every run) as we're changing behaviors to an preexisting script?
- Related to the above, should this version of the migration script check the API's version to ensure the new merge capabilities? In a real world scenario, a site admin may forget to upgrade all clusters before running this.
- In the story description, it is required a `--migrate-to` flag to mass assign a home cluster
- When using a tokens file, maybe it would be convenient to check for LoginCluster consistency as it's being done in the `envvar` case.

## #29 - 09/26/2019 07:40 PM - Lucas Di Pentima

Some more comments/questions:

- File `sdk/python/arvados/commands/federation_migrate.py`
  - The `main()` func has ~300 lines, could it be splitted up for enhanced readability/maintainability? I think it would also enable us to write some unit tests.
  - Lines 174-177 repeat at 202-205, could those be converted to a function?
  - At line ~176: what happens when there's a username conflict? If there're users A and B, and the idea is to rename A to B, B will be first renamed to 'Bmigrate', on the next pass when it's Bmigrate's turn, the script will see if should name the user as B, so former A (now B) user will be renamed to Bmigrate and former B will get its username back?
    - I did a test on this and the behavior is more or less what I was supposing, with some differences:  
Initial report:

```
email,username,user uuid,home cluster
ldipenti@gmail.com,ldipenti,x2uot-tpzed-5bq9jvuhxk0pepf,x2uot
ldipentima@veritasgenetics.com,ldipentima,x2uot-tpzed-icv1lofyc0dafs2,
ldipentima@veritasgenetics.com,ldipentima,x2yhh-tpzed-5u94ijf24fkeuiw,
ldipentima@veritasgenetics.com,ldipentima,xg6ao-tpzed-3xf9m0apph5m34m,
lucas@example.com,lucas,x2yhh-tpzed-x5r8hgvydks0q8e,x2yhh
test@example.com,test,x2yhh-tpzed-u47hwp50x1ab2js,x2yhh
```

Changed username 'ldipenti' to 'lucas' and set home cluster to 'x2yhh' and ran the script with `--migrate`:

```
$ arv-federation-migrate --tokens ~/tokens.csv --migrate /tmp/report.txt
Reading /home/lucas/tokens.csv
Contacting 172.17.0.2:8000
Contacting 172.17.0.3:8000
Contacting 172.17.0.4:8000
Checking that the federation is well connected
Getting user list from x2uot
Getting user list from xg6ao
Getting user list from x2yhh
(ldipenti@gmail.com) No user listed with same email to migrate x2uot-tpzed-5bq9jvuhxk0pepf
to x2yhh, will create new user with username 'lucas'
(ldipenti@gmail.com) Activating user x2yhh-tpzed-r4hm9zdybobmatr on x2uot
(ldipenti@gmail.com) Migrating x2uot-tpzed-5bq9jvuhxk0pepf to x2yhh-tpzed-r4hm9zdybobmatr on x2uot
(ldipentima@veritasgenetics.com) Skipping x2uot-tpzed-icv1lofyc0dafs2, no home cluster specified
(ldipentima@veritasgenetics.com) Skipping x2yhh-tpzed-5u94ijf24fkeuiw, no home cluster specified
(ldipentima@veritasgenetics.com) Skipping xg6ao-tpzed-3xf9m0apph5m34m, no home cluster specified
(lucas@example.com) Updating username of x2yhh-tpzed-x5r8hgvydks0q8e to 'lucas' on x2yhh
(lucas@example.com) Error updating username of x2yhh-tpzed-x5r8hgvydks0q8e
to 'lucas' on x2yhh: <HttpError 422 when requesting https://172.17.0.2:8000/arvados/v1/users/x2yhh-
tpzed-r4hm9zdybobmatr?alt=json returned "Username has already been taken (req-6yere4ls4xw10orzudgh
)">
```

After that, I ran the script to get a new report:

```
email,username,user uuid,home cluster
ldipenti@gmail.com,lucas,x2yhh-tpzed-r4hm9zdybobmatr,x2yhh
ldipentima@veritasgenetics.com,ldipentima,x2uot-tpzed-icvilofyc0dafs2,
ldipentima@veritasgenetics.com,ldipentima,x2yhh-tpzed-5u94ijf24fkeuiw,
ldipentima@veritasgenetics.com,ldipentima,xg6ao-tpzed-3xf9m0apph5m34m,
lucas@example.com,lucasmigrate,x2yhh-tpzed-x5r8hgvydks0g8e,x2yhh
test@example.com,test,x2yhh-tpzed-u47hwp50x1ab2js,x2yhh
```

The former 'lucas' user got renamed to 'lucasmigrate' and the script tried to rename it back to lucas, failing. Is this expected?

- The story description mentions testing with mocks. Were those dropped in favor of the integration tests with arvbox?
- The merge-remote-account doc page might need some updates

### #30 - 09/26/2019 07:45 PM - Peter Amstutz

- Description updated

Lucas Di Pentima wrote:

More comments & questions while I continue reviewing:

- Is it necessary to commit .cwl files that are generated from .cwllex?

I didn't want to require another thing to do just to be able to run tests.

- Typo at services/api/app/models/user.rb:L304

Fixed.

- "hotreset" documentation is missing on arvbox's help output

Fixed.

- File sdk/python/arvados/commands/federation\_migrate.py
  - At line 34 I got an IndexError exception when using a tokens file with an empty line (the last line), maybe it would be nice to ignore those.

Fixed.

- There's commented out code at line 60

Removed.

- Seems that there's repeated code blocks at lines 112-121 & 125-135.
  - I'm not sure what they do, maybe add some comments?
  - If the first for-loop is to get homeuuid to be empty string or some unique uuid on a list, wouldn't be something like the following be more clear?:  
[...]

Refactored & added comment.

- Do you think it is convenient to add a --version argument (and maybe report the version to stdout on every run) as we're changing behaviors to an preexisting script?

Added --version

- Related to the above, should this version of the migration script check the API's version to ensure the new merge capabilities? In a real world scenario, a site admin may forget to upgrade all clusters before running this.

Unfortunately I don't know of a good way to test that, the discovery document is the same since the main thing that changed is the ability to merge a remote user to a local user, which would fail prior to this branch. Ideas?

- In the story description, it is required a --migrate-to flag to mass assign a home cluster

I think the flag is redundant when you have LoginCluster set. Updated the description.

- When using a tokens file, maybe it would be convenient to check for LoginCluster consistency as it's being done in the envvar case.

I moved the code that determines LoginCluster out one level so it runs for both cases.

### #31 - 09/26/2019 08:32 PM - Peter Amstutz

Lucas Di Pentima wrote:

Some more comments/questions:

- File `sdk/python/arvados/commands/federation_migrate.py`
  - The `main()` func has ~300 lines, could it be splitted up for enhanced readability/maintainability? I think it would also enable us to write some unit tests.
  - Lines 174-177 repeat at 202-205, could those be converted to a function?

I'll see what refactoring makes sense.

- At line ~176: what happens when there's a username conflict? If there're users A and B, and the idea is to rename A to B, B will be first renamed to 'Bmigrate', on the next pass when it's Bmigrate's turn, the script will see if should name the user as B, so former A (now B) user will be renamed to Bmigrate and former B will get its username back?
  - I did a test on this and the behavior is more or less what I was supposing, with some differences:
    - Initial report:
      - [...]
      - Changed username 'ldipenti' to 'lucas' and set home cluster to 'x2yhh' and ran the script with `--migrate`:
      - [...]
      - After that, I ran the script to get a new report:
        - [...]
        - The former 'lucas' user got renamed to 'lucasmigrate' and the script tried to rename it back to lucas, failing. Is this expected?

If I understand correctly, you associated the "lucas" username with both [ldipenti@gmail.com](mailto:ldipenti@gmail.com) and [lucas@example.com](mailto:lucas@example.com) on x2yhh? So this a user error (you asked it to do conflicting things) but the script should be able to detect that and tell you that you made a mistake.

- The story description mentions testing with mocks. Were those dropped in favor of the integration tests with arvbox?

Yes. The things that I most needed to test were the specific API behaviors with user merging and handling usernames (which turned a bug and some hidden requirements). Mocks wouldn't have allowed me to discover those problems so I would still need to create the integration tests. I'll look at refactoring and see if it makes sense to add a few retrospective unit tests.

- The merge-remote-account doc page might need some updates

Will check it out.

### #32 - 09/30/2019 02:11 PM - Peter Amstutz

15531-logincluster-migrate @ [2602abaae640b1c50e9780d951b15006e43e203a](#)

- Refactor federation-migrate to break it into functions
- Check API server 'revision'
- Check for conflicting usernames
- Update documentation

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1561/>

### #33 - 09/30/2019 09:02 PM - Eric Biagiotti

Peter Amstutz wrote:

15531-logincluster-migrate @ [2602abaae640b1c50e9780d951b15006e43e203a](#)

- Refactor federation-migrate to break it into functions
- Check API server 'revision'
- Check for conflicting usernames
- Update documentation

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1561/>

- Does this affect WB/WB2 account merge functionality? Specifically wondering about the new admin remote to local case.
- `admin/merge-remote-account.html` - The first mention of LoginCluster could use a little explanation, maybe just a clause at the end of the sentence "as specified in the config". Also, its probably worth adding a sentence to the "Migrate users" section about using `--dry-run` before

running --migrate.

- fed-migrate/README - Does this need updating? arvbox-make-federation.cwl is complaining about missing --arvbox\_base.
- Do we have any tests that make sure object ownership is transferred correctly?

#### #34 - 10/04/2019 07:03 PM - Peter Amstutz

Eric Biagiotti wrote:

Peter Amstutz wrote:

15531-logincluster-migrate @ [2602abaae640b1c50e9780d951b15006e43e203a](#)

- Refactor federation-migrate to break it into functions
- Check API server 'revision'
- Check for conflicting usernames
- Update documentation

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1561/>

- Does this affect WB/WB2 account merge functionality? Specifically wondering about the new admin remote to local case.

I don't think so.

The current workbench UI doesn't allow linking a remote user to a local user, that hasn't changed.

Linking a remote user to a local one moves the remote user's stuff, but not login credentials.

If the remote user is admin and the local one isn't, the user would lose admin access. But that's the same as linking a local user (admin) to a (non-admin) remote one. The "check if old user is admin and new user is not" check is done by the client.

In the LoginCluster case, if a user is admin on the master, they are admin everywhere, so it only matters if the "new" user is admin.

- admin/merge-remote-account.html - The first mention of LoginCluster could use a little explanation, maybe just a clause at the end of the sentence "as specified in the config". Also, it's probably worth adding a sentence to the "Migrate users" section about using --dry-run before running --migrate.

Done.

- fed-migrate/README - Does this need updating? arvbox-make-federation.cwl is complaining about missing --arvbox\_base.

Done.

- Do we have any tests that make sure object ownership is transferred correctly?

api/test/integration/users\_test.rb includes object ownership checks.

#### #35 - 10/04/2019 07:04 PM - Peter Amstutz

15531-logincluster-migrate @ [a798952d64584cdcec2a13dc0df434692d2b521e](#)

#### #36 - 10/07/2019 06:58 PM - Eric Biagiotti

Peter Amstutz wrote:

15531-logincluster-migrate @ [a798952d64584cdcec2a13dc0df434692d2b521e](#)

- /admin/federation.html#LoginCluster - Might be useful to indicate what arvados services need to be restarted after adding LoginCluster. Similarly, the Configuration section above might benefit from this as well.
- When running cwltool fed-migrate.cwl fed.json, the workflow fails with the following.

```
INFO Could not collect memory usage, job ended before monitoring began.
INFO [job superuser_tok_3_3] completed success
INFO [step superuser_tok_3_3] completed success
INFO [workflow main_2_3] completed success
INFO [step main_2] completed success
INFO [workflow ] starting step run_test_3
INFO [step run_test_3] start
INFO [workflow run_test_3] start
```

```

INFO [workflow run_test_3] starting step main_1_2
INFO [step main_1_2] start
INFO [job main_1] /tmp/Ujy7Uo$ python \
  /tmp/tmpUj8wSe/stg4fe44f31-5c17-4e86-9d9a-1ec8c9d54572/create_users.py \
  _script
Traceback (most recent call last):
  File "/tmp/tmpUj8wSe/stg4fe44f31-5c17-4e86-9d9a-1ec8c9d54572/create_users.py", line 23, in <module>
    arvados.api(host=j["arvados_api_hosts"][1], token=maketoken(newtok), insecure=True).users().current().execute()
  File "/usr/local/lib/python2.7/dist-packages/google_api_python_client-1.6.7-py2.7.egg/googleapiclient/_helpers.py", line 130, in positional_wrapper
    return wrapped(*args, **kwargs)
  File "/usr/local/lib/python2.7/dist-packages/google_api_python_client-1.6.7-py2.7.egg/googleapiclient/http.py", line 840, in execute
    raise HttpError(resp, content, uri=self.uri)
arvados.errors.ApiError: <HttpError 401 when requesting https://172.17.0.3:8000/arvados/v1/users/current?alt=json returned "Not logged in (req-n4f34ddkc4hqfyehfcwu)">
INFO [job main_1] Max memory used: 48MiB
WARNING [job main_1] completed permanentFail
WARNING [step main_1_2] completed permanentFail
INFO [workflow run_test_3] completed permanentFail
WARNING [step run_test_3] completed permanentFail
INFO [workflow ] completed permanentFail
{
  "report": null,
  "supertok": [
    "5pkjz8u7ip4jrzedabr8jats0b5p482uuxp6k91p5tu41pa3e",
    "49sz6k2po2q6p66n9qu8ovzfx3om04ah195m7zsvf5hp5v9fiv",
    "m0ay7h4i9dvfpono3v3krqo7wz35sp7y26qjcfz93eipsrzb"
  ]
}

```

### #37 - 10/09/2019 01:25 PM - Peter Amstutz

Eric Biagiotti wrote:

Peter Amstutz wrote:

15531-logincluster-migrate @ [a798952d64584cdcec2a13dc0df434692d2b521e](#)

- /admin/federation.html#LoginCluster - Might be useful to indicate what arvados services need to be restarted after adding LoginCluster. Similarly, the Configuration section above might benefit from this as well.

Updated documentation

- When running cwltool fed-migrate.cwl fed.json, the workflow fails with the following.

[...]

Added switching git branches to the test setup, try it now. Also make sure to update your arvbox docker image (arvbox update).

### #38 - 10/09/2019 03:07 PM - Peter Amstutz

- Target version changed from 2019-10-09 Sprint to 2019-10-23 Sprint

### #39 - 10/09/2019 08:24 PM - Eric Biagiotti

Peter Amstutz wrote:

Eric Biagiotti wrote:

Peter Amstutz wrote:

15531-logincluster-migrate @ [a798952d64584cdcec2a13dc0df434692d2b521e](#)

- /admin/federation.html#LoginCluster - Might be useful to indicate what arvados services need to be restarted after adding LoginCluster. Similarly, the Configuration section above might benefit from this as well.

Updated documentation

- When running `cwltool fed-migrate.cwl fed.json`, the workflow fails with the following.

[...]

Added switching git branches to the test setup, try it now. Also make sure to update your arvbox docker image (arvbox update).

Thanks, tests ran successfully. A few more things:

- Running `cwltool fed-migrate.cwl fed.json` with a py2.7 version of `arv-federation-migrate` fails with the following error. Not sure if this is a priority for us, but we should at least note that its py3 only in the docs.

```
(case3@test) No user listed with same email to migrate x371k-tpzed-195x5allw5qt33k
to xcwah, will create new user with username 'case3'
Traceback (most recent call last):
  File "/home/eric/envs/fed2/bin/arv-federation-migrate", line 7, in <module>
    exec(compile(f.read(), __file__, 'exec'))
  File "/home/eric/projects/arvados/sdk/python/bin/arv-federation-migrate", line 7, in <module>
    main()
  File "/home/eric/projects/arvados/sdk/python/arvados/commands/federation_migrate.py", line 358, in main
    newuser = activate_remote_user(args, email, homearv, migratearv, old_user_uuid, new_user_uuid)
  File "/home/eric/projects/arvados/sdk/python/arvados/commands/federation_migrate.py", line 241, in activate_
remote_user
    digestmod='sha1').hexdigest()
  File "/usr/lib/python2.7/hmac.py", line 136, in new
    return HMAC(key, msg, digestmod)
  File "/usr/lib/python2.7/hmac.py", line 52, in __init__
    self.outer = self.digest_cons()
  File "/usr/lib/python2.7/hmac.py", line 50, in <lambda>
    self.digest_cons = lambda d='': digestmod.new(d)
AttributeError: 'str' object has no attribute 'new'
```

- Looks like in `federation_migrate.py` line 259, we try and activate the new user if its not already active. What if the old user was also inactive, do we still want to try and activate the new user? Along these lines, should we be testing use cases where `arv-federation-migrate` will fail/print an error?

**#40 - 10/10/2019 10:07 PM - Peter Amstutz**

- Related to Story #15720: [API] Unified user listing across all clusters in a federation added

**#41 - 10/10/2019 10:07 PM - Peter Amstutz**

Thanks, tests ran successfully. A few more things:

- Running `cwltool fed-migrate.cwl fed.json` with a py2.7 version of `arv-federation-migrate` fails with the following error. Not sure if this is a priority for us, but we should at least note that its py3 only in the docs.

Fixed on Python 2.7. I expect we're supporting Python 2.7 for at least the next release.

- Looks like in `federation_migrate.py` line 259, we try and activate the new user if its not already active. What if the old user was also inactive, do we still want to try and activate the new user?

Good point. Now it won't activate the 'new' user if the 'old' user isn't active.

Along these lines, should we be testing use cases where `arv-federation-migrate` will fail/print an error?

In principal, yes, it's good to test error cases, but I feel like in this case it would be significantly more work for diminishing returns.

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1582/>

**#42 - 10/11/2019 02:34 PM - Eric Biagiotti**

Peter Amstutz wrote:

Thanks, tests ran successfully. A few more things:

- Running `cwltool fed-migrate.cwl fed.json` with a py2.7 version of `arv-federation-migrate` fails with the following error. Not sure if this is a priority for us, but we should at least note that its py3 only in the docs.

Fixed on Python 2.7. I expect we're supporting Python 2.7 for at least the next release.

- Looks like in `federation_migrate.py` line 259, we try and activate the new user if its not already active. What if the old user was also inactive, do we still want to try and activate the new user?

Good point. Now it won't activate the 'new' user if the 'old' user isn't active.

Along these lines, should we be testing use cases where `arv-federation-migrate` will fail/print an error?

In principal, yes, it's good to test error cases, but I feel like in this case it would be significantly more work for diminishing returns.

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1582/>

Thanks! Tests ran successfully in py2/3. This LGTM!

**#43 - 10/11/2019 06:55 PM - Peter Amstutz**

- *Status changed from In Progress to Resolved*

**#44 - 01/21/2020 09:20 PM - Peter Amstutz**

- *Release set to 22*