# Arvados - Feature #14718

## [API] Option to issue salted token in login procedure

01/10/2019 08:00 PM - Peter Amstutz

| | | | |
|---|---|---|---|
| **Status:** | Resolved | **Start date:** | 01/26/2019 |
| **Priority:** | Normal | **Due date:** | |
| **Assigned To:** | Lucas Di Pentima | **% Done:** | 100% |
| **Category:** | API | **Estimated time:** | 0.00 hour |
| **Target version:** | 2019-01-30 Sprint | | |

**Description**

When redirecting a user agent to the Arvados login endpoint (https://aaaaa.arvadosapi.com/login), an application may specify a remote cluster ID in the query string (.../login?remote=bbbbb).

In that case, assuming the login is successful and a new token is issued, the new token will be salted for the specified remote.

This story does *not* cover prompting the user differently depending on the requested scope, although we will want to do so in the future.

**Subtasks:**

| | |
|---|---|
| Task # 14751: Review 14718-api-login-salted-token | **Resolved** |

**Related issues:**

| | |
|---|---|
| Blocks Arvados - Feature #12958: [Federation] Workbench login chooser | **Closed** |

### Associated revisions

**Revision d1e00d89 - 01/29/2019 09:30 AM - Lucas Di Pentima**

Merge branch '14718-api-login-salted-token'
Closes #14718

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

### History

**#1 - 01/10/2019 08:00 PM - Peter Amstutz**

*- Status changed from New to In Progress*

**#2 - 01/10/2019 08:14 PM - Peter Amstutz**

*- Subject changed from [API] untrusted client login receives salted token to [API] login flow for remote clusters that receives salted token*

**#4 - 01/10/2019 08:42 PM - Peter Amstutz**

*- Related to Feature #12958: [Federation] Workbench login chooser added*

**#5 - 01/10/2019 08:46 PM - Tom Clegg**

*- Subject changed from [API] login flow for remote clusters that receives salted token to [API] Option to issue salted token in login procedure*

*- Description updated*

*- Category set to API*

*- Status changed from In Progress to New*

**#6 - 01/10/2019 08:46 PM - Tom Clegg**

*- Related to deleted (Feature #12958: [Federation] Workbench login chooser)*

**#7 - 01/10/2019 08:46 PM - Tom Clegg**

*- Blocks Feature #12958: [Federation] Workbench login chooser added*

**#10 - 01/10/2019 08:47 PM - Tom Clegg**

*- Tracker changed from Bug to Feature*

#### #11 - 01/10/2019 09:30 PM - Tom Clegg

*- Story points set to 2.0*

#### #12 - 01/11/2019 07:13 PM - Peter Amstutz

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

#### #13 - 01/16/2019 04:15 PM - Tom Morris

*- Target version changed from Arvados Future Sprints to 2019-01-30 Sprint*

#### #14 - 01/16/2019 04:42 PM - Lucas Di Pentima

*- Assigned To set to Lucas Di Pentima*

#### #15 - 01/18/2019 04:11 PM - Lucas Di Pentima

*- Status changed from New to In Progress*

#### #16 - 01/26/2019 12:11 AM - Lucas Di Pentima

Updates at 4bb449eb5 - branch 14718-api-login-salted-token
Test run: https://ci.curoverse.com/job/developer-run-tests/1041/

Took advantage of return_to being passed around between API server and SSO to propagate the remote parameter.

#### #17 - 01/28/2019 04:18 PM - Tom Clegg

In user_sessions_controller.rb → login, it looks like params[:remote] is missing a CGI.escape() in case it's an unexpectedly non-alphanum/malicious string -- it gets appended to another string provided by the same client so I don't see how it would be exploitable, but I'm inclined to be defensive about it anyway.

```
        remote_param += "remote=#{params[:remote]}"
```

Stashing the remote param in the return_to URL seems sensible enough, but

- it's sneaky/non-obvious enough that it should be explained in comments
- it should be removed from return_to before create() redirects there. If I'm following correctly, when Workbench sends the user to https://api/login?return_to=https://wb/home&remote=bbbbb, the user will eventually land on either https://wb/home?remote=bbbbb or https://wb/home depending on whether they needed to go through the joshid→sso→sessions#create flow or took the "already logged in" short cut. It would be better if we could ensure they always land on https://wb/home as requested. (Am I following correctly?)

I wonder whether it might even be less confusing to use a format like return_to="bbbbb,https://wb/home" (or ",https://wb/home" for the unsalted case) to avoid all the corner cases involved in stashing the remote id in the client-provided URL. For example, I'd rather not even try to figure out whether we're doing the right thing here when the client-provided return_to is something like "https://wb/home?foo=bar#foo/bar?baz".

#### #18 - 01/28/2019 08:17 PM - Lucas Di Pentima

Updates at 0d1836a8d
Test run: https://ci.curoverse.com/job/developer-run-tests/1043/

Addressed above suggestions.

#### #19 - 01/28/2019 08:53 PM - Tom Clegg

I think we should

- ignore the remote param if it contains a comma (or perhaps even if it doesn't match /^[0-9a-z]{5}$/)
- escape the whole combined param including the comma (I notice CGI.escape(',') == '%2c' -- and we will unescape the whole thing before splitting, so we might as well escape the whole thing after joining)

the rest LGTM, thanks!

#### #20 - 01/28/2019 08:55 PM - Tom Clegg

Tom Clegg wrote:

> - ignore the remote param if it contains a comma (or perhaps even if it doesn't match /^[0-9a-z]{5}$/)

...only on the "login" handler, not the "callback" handler.

#### #21 - 01/28/2019 09:26 PM - Tom Clegg

On second thought, I think we should error out (rather than ignoring the remote param) in both callback and login handlers if the "remote" value

doesn't look like a cluster ID.

#### #22 - 01/28/2019 10:06 PM - Lucas Di Pentima

Updates at 245e4cafd
Test run: https://ci.curoverse.com/job/developer-run-tests/1044/

- Validates the remote cluster id parameter both on login endpoint & omniauth callback.
- Added tests.
- Updated lib/controller handler test to support the new format.

#### #23 - 01/28/2019 10:32 PM - Tom Clegg

LGTM, thanks!

#### #24 - 01/29/2019 12:40 PM - Lucas Di Pentima

*- Status changed from In Progress to Resolved*

*- % Done changed from 0 to 100*

Applied in changeset arvados|d1e00d89dac87929d39c0689a593f0574980f2e8.

#### #25 - 03/01/2019 06:30 PM - Tom Morris

*- Release set to 15*