

Arvados - Bug #17803

[config-check] failing to warn about unrecognized config keys

06/11/2021 06:23 PM - Tom Clegg

Status:	Resolved	Start date:	06/13/2021
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2021-06-23 sprint		
Description			
It seems that the "warn about unrecognized config key" feature does not warn about keys that are unrecognized because of upper/lowercase differences -- so it warns about BadKey (as in the test case), but not RunTimeEngine (where there is a real config key RuntimeEngine).			
Loading and bad key detection should both be case-sensitive, or both be case-insensitive.			
Subtasks:			
Task # 17804: Review 17803-config-key-case-warning			Resolved

Associated revisions

Revision 9f06f5b2 - 06/16/2021 01:45 PM - Tom Clegg

Merge branch '17803-config-key-case-warning'

fixes #17803

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tom@curii.com>

Revision 7c9ae0db - 07/09/2021 02:39 PM - Ward Vandewege

Small lib/config load_test.go change to reflect the absence of the RuntimeEngine config variable in the 2.2 branch.

refs #17803

Arvados-DCO-1.1-Signed-off-by: Ward Vandewege <ward@curii.com>

History

#1 - 06/13/2021 02:08 PM - Tom Clegg

- Release set to 39
- Target version set to 2021-06-23 sprint
- Assigned To set to Tom Clegg
- Status changed from New to In Progress

17803-config-key-case-warning @ [4e65a41bea4b892ef7232bfe6b9b20ca35380368](https://ci.arvados.org/view/Developer/job/developer-run-tests/2520/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2520/>

Warning with this fix:

```
deprecated or unknown config entry: Clusters.zzzzz.Containers.RunTimeEngine (perhaps you meant RuntimeEngine?)
```

#2 - 06/14/2021 05:12 PM - Ward Vandewege

Tom Clegg wrote:

17803-config-key-case-warning @ [4e65a41bea4b892ef7232bfe6b9b20ca35380368](https://ci.arvados.org/view/Developer/job/developer-run-tests/2520/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2520/>

Warning with this fix:

[...]

The patch works, thanks. One question, consider this config file:

Clusters:

```
xxxxx:
  # Token used internally by Arvados components to authenticate to
  # one another. Use a string of at least 50 random alphanumerics.
  SystemrootToken: ""
```

Services:

```
# In each of the service sections below, the keys under
# InternalURLs are the endpoints where the service should be
# listening, and reachable from other hosts in the cluster.
KeepStore:
  Internalurls:
    "http://host1.example:12345": {}
    "http://host2.example:12345":
      # Rendezvous is normally empty/omitted. When changing the
      # URL of a Keepstore service, Rendezvous should be set to
      # the old URL (with trailing slash omitted) to preserve
      # rendezvous ordering.
      RendezVous: "xxxx"
  ExternalURL: "-"
```

Containers:

```
RunTimeEngine: "xxxx"
```

I've deliberately screwed up the case of `SystemrootToken`, `ServiceS`, `KeepStore`, `Internalurls`, `RendezVous`, and `RunTimeEngine`.

This the output from `config-check`:

```
time="2021-06-14T13:01:30-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.Containers.RunTimeEngine (perhaps you meant RuntimeEngine?)"
time="2021-06-14T13:01:30-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.ServiceS (perhaps you meant Services?)"
time="2021-06-14T13:01:30-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.SystemrootToken (perhaps you meant SystemRootToken?)"
time="2021-06-14T13:01:30-04:00" level=warning msg="Clusters.xxxxx.ManagementToken: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:01:30-04:00" level=warning msg="Clusters.xxxxx.SystemRootToken: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:01:30-04:00" level=warning msg="Clusters.xxxxx.Collections.BlobSigningKey: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:01:30-04:00" level=warning msg="keepstore configured at http://host1.example:12345/ does not have access to any volumes"
time="2021-06-14T13:01:30-04:00" level=warning msg="keepstore configured at http://host2.example:12345/ does not have access to any volumes"
```

It complains about `RunTimeEngine`, `ServiceS`, `SystemrootToken`. It does not see the mistake in `KeepStore`, `Internalurls` and `RendezVous`, because it ignored the `ServiceS` section.

But... it is clearly still parsing the `ServiceS` block, as it complains about the defined keepstores not having access to any volumes. What's up with that?

Another example, this one has the spelling of `Services` fixed:

Clusters:

```
xxxxx:
  # Token used internally by Arvados components to authenticate to
  # one another. Use a string of at least 50 random alphanumerics.
  SystemrootToken: ""
```

Services:

```
# In each of the service sections below, the keys under
# InternalURLs are the endpoints where the service should be
# listening, and reachable from other hosts in the cluster.
KeepStore:
  Internalurls:
    "http://host1.example:12345": {}
    "http://host2.example:12345":
      # Rendezvous is normally empty/omitted. When changing the
      # URL of a Keepstore service, Rendezvous should be set to
      # the old URL (with trailing slash omitted) to preserve
      # rendezvous ordering.
      RendezVous: "xxxx"
  ExternalURL: "-"
```

Containers:

```
RunTimeEngine: "xxxx"
```

That leads to this output:

```
time="2021-06-14T13:09:59-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.Containers.RunTimeEngine (perhaps you meant RuntimeEngine?)"
time="2021-06-14T13:09:59-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.Services.KeepStore.Internalurls (perhaps you meant InternalURLs?)"
time="2021-06-14T13:09:59-04:00" level=warning msg="deprecated or unknown config entry: Clusters.xxxxx.SystemRootToken (perhaps you meant SystemRootToken?)"
time="2021-06-14T13:09:59-04:00" level=warning msg="Clusters.xxxxx.ManagementToken: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:09:59-04:00" level=warning msg="Clusters.xxxxx.SystemRootToken: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:09:59-04:00" level=warning msg="Clusters.xxxxx.Collections.BlobSigningKey: secret token is not set (use 32+ random characters from a-z, A-Z, 0-9)"
time="2021-06-14T13:09:59-04:00" level=warning msg="keepstore configured at http://host1.example:12345/ does not have access to any volumes"
time="2021-06-14T13:09:59-04:00" level=warning msg="keepstore configured at http://host2.example:12345/ does not have access to any volumes"
```

In this case, it ignores the case difference in `KeepStore`, but it complains about `Internalurls`. And it still complains about the access to volumes.

#3 - 06/14/2021 05:30 PM - Tom Clegg

Ah. The "merge actual into defaults" operation, given differing-case keys with different map values, merges the two maps (because unmarshal into a non-empty map adds/updates but doesn't clear first).

<https://play.golang.org/p/0EkcL5ii5WI>

I think we can fix the strangeness by deleting the unrecognized keys as well as reporting them.

#4 - 06/14/2021 06:39 PM - Tom Clegg

Added ServiceS, KeepStore, and RendezVous to the test case and made the necessary adjustments so they all get reported.

I noticed a couple of remaining oddities:

- We could warn about using Rendezvous in a Services entry other than Keepstore (but we don't).
- We should reject inputs with duplicate keys (but we don't). It looks like upgrading to gopkg.in/yaml.v3 would help, but that means upgrading ghodss/yaml to use it, and... at this point it starts to feel out-of-scope.

17803-config-key-case-warning @ [4b5ec8651ba83f2a79fe40708021e03c86275093](https://ci.arvados.org/view/Developer/job/developer-run-tests/2524/) --
<https://ci.arvados.org/view/Developer/job/developer-run-tests/2524/>

#5 - 06/15/2021 08:11 PM - Ward Vandewege

Tom Clegg wrote:

Added ServiceS, KeepStore, and RendezVous to the test case and made the necessary adjustments so they all get reported.

I noticed a couple of remaining oddities:

- We could warn about using Rendezvous in a Services entry other than Keepstore (but we don't).
- We should reject inputs with duplicate keys (but we don't). It looks like upgrading to gopkg.in/yaml.v3 would help, but that means upgrading ghodss/yaml to use it, and... at this point it starts to feel out-of-scope.

17803-config-key-case-warning @ [4b5ec8651ba83f2a79fe40708021e03c86275093](https://ci.arvados.org/view/Developer/job/developer-run-tests/2524/) --
<https://ci.arvados.org/view/Developer/job/developer-run-tests/2524/>

This works; the only thing I don't like is the use of an invalid config (rendezvous on Keepproxy record) in the tests. Can we not define a volume in the mock to make the "keepstore has no volumes" warning go away instead? It seems awkward to do it this way, particularly because there's a comment in the mock that Rendezvous on keepproxy should be another warning.

Otherwise, LGTM. Should we list the known limitations you mentioned above as FIXME's in `lib/config/load.go`? I like how e.g. the tailscale codebase is full of those kinds of comments. It seems more easily actionable to record them that way (as opposed to creating tickets we may not look at for years).

#6 - 06/15/2021 10:32 PM - Tom Clegg

define a volume in the mock to make the "keepstore has no volumes" warning go away

Oh, yeah, that's a better solution! Done.

FIXME's

Added.

17803-config-key-case-warning @ [2d3723c301085ebfcb1f3a8940451edc15f10f93](https://ci.arvados.org/view/Developer/job/developer-run-tests/2529/) --
<https://ci.arvados.org/view/Developer/job/developer-run-tests/2529/>

#7 - 06/16/2021 01:19 PM - Ward Vandewege

Tom Clegg wrote:

define a volume in the mock to make the "keepstore has no volumes" warning go away

Oh, yeah, that's a better solution! Done.

FIXME's

Added.

17803-config-key-case-warning @ [2d3723c301085ebfcb1f3a8940451edc15f10f93](https://ci.arvados.org/view/Developer/job/developer-run-tests/2529/) --
<https://ci.arvados.org/view/Developer/job/developer-run-tests/2529/>

LGTM, thanks!

#8 - 06/16/2021 02:47 PM - Tom Clegg

- % Done changed from 0 to 100

- Status changed from In Progress to Resolved

Applied in changeset [arvados!9f06f5b2e5f5549c61458db7f401fc2e2dc7dd0b](https://ci.arvados.org/view/Developer/job/developer-run-tests/2529/).