

Arvados - Feature #17607

[deployment][provision][arvados-formula] register shell nodes with Arvados

04/29/2021 01:44 PM - Javier Bértoli

Status:	Resolved	Start date:	05/17/2021
Priority:	Normal	Due date:	
Assigned To:	Javier Bértoli	% Done:	100%
Category:	Deployment	Estimated time:	0.00 hour
Target version:	2021-07-07 sprint		
Description			
When adding shell nodes, we want to register them to Arvados at creation time. This can be done running			
<pre>arv --format=uuid virtual_machine create --virtual-machine '{"hostname":"shell.ClusterID.example.com"}'</pre>			
using the SystemRootToken.			
When the shell node finishes deployment, the command above should be run (verifying first that it does not already exist)			
As this is an optional step that should be decided by the cluster Admin, we should only add this as an example.			
Also, this can be done in different ways (ie, in a master/minion salt deployment, using the mine).			
We can add an extra example state in the formula, which can then be added in the shell role by the provision script.			
Subtasks:			
Task # 17640: Review #17607 [deployment][provision][arvados-formula] register shell nod...			Resolved
Related issues:			
Blocks Arvados - Story #17512: Release Arvados 2.2		Resolved	05/03/2021

History

#1 - 05/04/2021 07:25 PM - Peter Amstutz

- Blocks Story #17512: Release Arvados 2.2 added

#2 - 05/05/2021 03:54 PM - Javier Bértoli

- Blocks Bug #17601: [deployment][arvados-formula] webshell needs a cron for login-sync added

#3 - 05/05/2021 03:56 PM - Javier Bértoli

- Status changed from New to In Progress

#4 - 05/10/2021 03:46 PM - Peter Amstutz

Related to shell nodes, it would be extremely useful if sudo could be set up to support the same PAM config as webshell, so that users can use valid Arvados tokens to authenticate to sudo.

#5 - 05/12/2021 03:52 PM - Peter Amstutz

- Target version changed from 2021-05-12 sprint to 2021-05-26 sprint

#6 - 05/13/2021 03:28 PM - Peter Amstutz

- Release set to 38

#7 - 05/14/2021 10:23 PM - Javier Bértoli

Added a new state to manage **Arvados resources** through the formula.

These resources can be configured through the pillar, under `arvados.cluster.resources.virtual_machines`

Initially, added an `arvados.controller.resources.virtual_machines` state file, to manage virtual machines and their scoped token as described in the documentation.

More resources can be added in the future easily.

The formula creates an entry for the virtual machine if it does not exist, checks for existing scoped tokens for it and adds one if none is found.

As both virtual machines and scoped tokens for them can be created by other means and Arvados does not check for their uniqueness, the formula uses the first existing entry for both resources if more than one are found.

I added checks to avoid creating these resources if they're found for a given machine.

The **nginx** configuration examples for these resources is still missing.

commit 1c81e52@arvados-formula, branch 17607-register-shell-nodes

#8 - 05/17/2021 06:30 PM - Javier Bértoli

Added nginx examples to add multiple webshells (commit 5bf76f7@arvados-formula, 17607-register-shell-nodes)

#9 - 05/17/2021 10:14 PM - Javier Bértoli

Added an example of some crude orchestration to get a shell node's scoped token and vm uuid and add them in a cron for arvados-login-sync.

As these tasks require *orchestration* (api server must be running before these queries being possible), I added them as examples, as they're useless on a single node configuration (which is the scope of a salt formula)

commit 4040230@arvados-formula, branch 17607-register-shell-nodes.

Need testing

#10 - 05/17/2021 10:19 PM - Javier Bértoli

- Status changed from *In Progress* to *Feedback*

#11 - 05/18/2021 06:52 PM - Lucas Di Pentima

Some comments:

- File arvados/controller/resources/virtual_machines.sls
 - Line 45: Should 'fixme' be on the regex? AFAICT this grep would always fail, or am I missing something? I'm seeing fixme being used on tests, but this grep command also would get executed on minions, correct? Shouldn't that regex use something like {a-z0-9}[5] instead?
- File arvados/controller/service/running.sls
 - Line 34-36: Is the unless clause necessary provided that it's using the same command execution check that que command itself?
- File test/salt/states/examples/arvados/shell/cron/add-login-sync.sls
 - Lines 58, 71, 81: Some refer to a file called /tmp/sc-tk<vm> and other to a file called /tmp/sctk-<vm> Are those references of the same file? If so, do you think it would be better to set the file name/path into a variable and use that?
 - Related to the above: Do you think it is feasible to randomize the use of temp files to avoid har dot debug issues?

#12 - 05/18/2021 07:43 PM - Javier Bértoli

Lucas Di Pentima wrote:

Some comments:

- File arvados/controller/resources/virtual_machines.sls
 - Line 45: Should 'fixme' be on the regex? AFAICT this grep would always fail, or am I missing something? I'm seeing fixme being used on tests, but this grep command also would get executed on minions, correct? Shouldn't that regex use something like {a-z0-9}[5] instead?

You're right. Fixed with a regex.

- File arvados/controller/service/running.sls
 - Line 34-36: Is the unless clause necessary provided that it's using the same command execution check that que command itself?

unless and onlyif are control parameters in saltstack, that prevent/ensure a state to run if the check is true.

In this case, the unless provide a 'cosmetic' functionality: If the unless check passes, Salt won't 'run the state' and it won't be shown in the run report. Otherwise, it will be reported every time the state is run, which might be confusing.

- File test/salt/states/examples/arvados/shell/cron/add-login-sync.sls
 - Lines 58, 71, 81: Some refer to a file called /tmp/sc-tk<vm> and other to a file called /tmp/sctk-<vm> Are those references of the same file? If so, do you think it would be better to set the file name/path into a variable and use that?

Right, fixed the name to be consistent, thanks.

- Related to the above: Do you think it is feasible to randomize the use of temp files to avoid har dot debug issues?

I don't wanted it to be a random name. In fact, I want it to be the same name: If the file is present, I don't want the state to run the command again. If we randomize it, every time the check will be run. And as we have to loop over all the scoped tokens, it's a costly run.

Squashed and force-pushed commit:79275e49|arvados-formula, branch 17607-register-shell-nodes

#13 - 05/18/2021 08:12 PM - Lucas Di Pentima

Just one additional comment:

- In file test/salt/states/examples/arvados/shell/cron/add-login-sync.sls line 71 there's one more sctk-<vm> file reference that seems to need correction. Note the dash '-', that isn't present on the other instances.

Apart from that, it LGTM.

#14 - 05/18/2021 08:18 PM - Javier Bértoli

Fixed at commit:e6dace0|arvados-formula , branch 17607-register-shell-nodes

#15 - 05/26/2021 04:10 PM - Peter Amstutz

- Target version changed from 2021-05-26 sprint to 2021-06-09 sprint

#16 - 05/28/2021 07:27 PM - Javier Bértoli

- Status changed from Feedback to In Progress

#17 - 05/28/2021 07:38 PM - Javier Bértoli

- Blocks deleted (Bug #17601: [deployment][arvados-formula] webshell needs a cron for login-sync)

#18 - 05/28/2021 10:51 PM - Lucas Di Pentima

LGTM

#19 - 06/09/2021 03:45 PM - Peter Amstutz

- Target version changed from 2021-06-09 sprint to 2021-06-23 sprint

#20 - 06/23/2021 03:34 PM - Peter Amstutz

- Target version changed from 2021-06-23 sprint to 2021-07-07 sprint

#21 - 06/23/2021 08:47 PM - Javier Bértoli

- Status changed from In Progress to Resolved

Merged