

openQA Infrastructure - action #89815

osd-deployment blocked by openqaworker-arm-3 offline and not recovered automatically

2021-03-10 08:35 - okurz

Status:	Resolved	Start date:	2021-03-10
Priority:	High	Due date:	2021-04-22
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
Observation			
<p>https://gitlab.suse.de/openqa/osd-deployment/pipelines shows that deployment on 2021-03-08 and 2021-03-10 have been blocked by openqaworker-arm-3 being offline. https://monitor.qa.suse.de/d/1bNU0StZz/automatic-actions?orgId=1 shows that openqaworker-arm-3 is offline since 2021-03-06 and automatic actions could not recover</p>			
Acceptance criteria			
<ul style="list-style-type: none">• AC1: openqaworker-arm-3 is online again• AC2: openqaworker-arm-3 can be recovered automatically or ticket to EngInfra is created automatically as in before• AC3: osd-deployment is recovered			
Suggestions			
<ul style="list-style-type: none">• Remove openqaworker-arm-3 from salt keys• Trigger osd-deployment again manually• Check automatic actions for 2021-03-06 and find out if a ticket was created or not, fix if not• Bring back openqaworker-arm-3 after the above is fixed (or get rid of this unreliable machine and replace it with a better one)			
Related issues:			
Related to openQA Infrastructure - action #92176: [alert] openqaworker-arm-3 ...		Resolved	2021-05-05 2021-05-21
Related to openQA Infrastructure - action #94456: no data from any arm host o...		Resolved	2021-06-22

History

#1 - 2021-03-10 08:35 - okurz

- Status changed from Workable to In Progress

- Assignee set to okurz

#2 - 2021-03-10 08:37 - okurz

Deleted salt key for openqaworker-arm-3 with `sudo salt-key -d openqaworker-arm-3.suse.de` and retriggered osd deployment pipeline on <https://gitlab.suse.de/openqa/osd-deployment/pipelines>. Running as <https://gitlab.suse.de/openqa/osd-deployment/-/pipelines/106382>

#3 - 2021-03-10 09:04 - okurz

- Status changed from In Progress to Workable

- Assignee deleted (okurz)

- Priority changed from Urgent to High

osd-deployment passed so far. As I have the salt key for the worker removed I consider this less critical now.

#4 - 2021-04-07 11:02 - mkittler

- Assignee set to mkittler

I'll have a look because with only 2 workers we can easily get into a situation where arm jobs aren't processed fast enough, see [#90755](#).

#5 - 2021-04-07 11:46 - mkittler

sol, power reset and power cycle don't work so it is not a surprise that our automatic recovery didn't work either. Maybe I can play around with IPMI a

little bit more. Otherwise I suppose we'll have to resort to filing an infra ticket.

#6 - 2021-04-07 11:55 - mkittler

Skip to the end if you're not interested in my exploratory use of ipmitool :-)

```
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis selftest
Self Test Results      : passed
```

```
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis restart_cause
System restart cause: chassis power control command
```

```
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis policy always-on
Set chassis power restore policy to always-on
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis power reset
Set Chassis Power Control to Reset failed: Command not supported in present state
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis power cycle
Set Chassis Power Control to Cycle failed: Command not supported in present state
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis power diag
Set Chassis Power Control to Diag failed: Command not supported in present state
```

```
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis power on
Chassis Power Control: Up/On
ipmitool -I lanplus -C 3 -H openqaworker-arm-3-ipmi.suse.de -U ADMIN -P ADMIN chassis power diag
Chassis Power Control: Diag
```

And something finally shows up on the SOL. It took a while as it conducted a few tests but now the system is booting. It looks like the chassis was turned off (whatever that means). Maybe our automatic recovery should also ensure that the chassis is turned on.

#7 - 2021-04-07 12:24 - okurz

Wow, that's interesting. Can we just add a chassis on in the grafana-webhook actions repo?

#8 - 2021-04-07 14:20 - mkittler

If "grafana-webhook" is the place where we do the automatic recovery, than that's what I'd do. Still have to find it, though.

#9 - 2021-04-07 14:29 - mkittler

- Status changed from Workable to In Progress

SR: https://gitlab.suse.de/openqa/grafana-webhook-actions/-/merge_requests/13

By the way, I've also added the worker back to salt. I did *not* resume the host-up alert because these seem to be generally disabled for the unstable arm workers.

#10 - 2021-04-08 04:16 - openqa_review

- Due date set to 2021-04-22

Setting due date based on mean cycle time of SUSE QE Tools

#11 - 2021-04-08 11:25 - mkittler

The SR has been merged and I've tested whether it works. It does. After power off the Grafana alert triggers the GitLab pipeline which then turns the power on again. However, since the power state is apparently not immediately reflected the subsequent power cycle command doesn't work:

```
PING openqaworker-arm-3.suse.de (10.160.0.85) 56(84) bytes of data.
--- openqaworker-arm-3.suse.de ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms
Chassis Power Control: Up/On
Set Chassis Power Control to Cycle failed: Command not supported in present state
```

Not a big deal because it isn't required anyways in that state. If the power was on and the worker just otherwise stuck the power on command will be a noop which zero return code so this case should be fine as well.

#12 - 2021-04-08 11:25 - mkittler

- Status changed from In Progress to Feedback

#13 - 2021-04-09 10:51 - okurz

discussed in weekly 2021-04-09:

<https://gitlab.suse.de/openqa/grafana-webhook-actions/-/blob/master/ipmi-health-check#L22> wastes 10 minutes and also we found out that

https://gitlab.suse.de/openqa/grafana-webhook-actions/-/merge_requests/12 did more than what we expected. I read all comments from [#76876](#) again but have not found any discussion that we want to turn the "do-reboot" steps into a "check-health" script. So mkittler and me decided that instead we want to go back to the previous way to trigger a reboot as soon as the script is started, not check if the machine might be still up (respond to ping but potentially not being fully reactive) and not waste 10 minutes of waiting if ping comes back successful before triggering any action.

Also I suggest:

```
if $ipmitool power status | grep -q 'is on'; then
    $ipmitool power cycle
else
    $ipmitool power on
fi
```

#14 - 2021-04-09 10:51 - mkittler

SR for improving the problem mentioned in the previous comment: https://gitlab.suse.de/openqa/grafana-webhook-actions/-/merge_requests/14 for AC2: https://gitlab.suse.de/openqa/grafana-webhook-actions/-/merge_requests/15

#15 - 2021-04-12 07:37 - okurz

- Status changed from Feedback to Resolved

Both MRs merged and considered stable. We won't test the actual ticket creation

#16 - 2021-05-05 14:08 - okurz

- Related to action #92176: [alert] openqaworker-arm-3 offline and CI pipeline unable to send email but stating "passed" added

#17 - 2021-06-22 22:00 - okurz

- Related to action #94456: no data from any arm host on <https://stats.openqa-monitor.qa.suse.de/d/1bNU0StZz/automatic-actions?orgId=1> added