

openQA Infrastructure - action #92969

Failing service os-autoinst-openvswitch after boot of some workers

2021-05-23 07:54 - okurz

Status:	Resolved	Start date:	2021-05-23
Priority:	High	Due date:	
Assignee:	okurz	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
Observation			
https://stats.openqa-monitor.qa.suse.de/d/KToPYLEWz/failed-systemd-services shows			
Currently failing services			
Last update Host Failing units # failed services			
2021-05-23 05:48:00 openqaworker-arm-1 var-lib-openqa-share.mount, os-autoinst-openvswitch 2			
2021-05-23 03:40:00 openqaworker13 var-lib-openqa-share.mount 1			
2021-05-23 03:39:00 grenache-1 var-lib-openqa-share.mount, os-autoinst-openvswitch 2			
2021-05-22 04:13:00 openqaworker-arm-2 var-lib-openqa-share.mount 1			
2021-05-22 01:47:00 openqaworker-arm-3 var-lib-openqa-share.mount, os-autoinst-openvswitch 2			
Acceptance criteria			
<ul style="list-style-type: none">• AC1: No failing os-autoinst-openvswitch after multiple reboot of many machines			
Suggestions			
<ul style="list-style-type: none">• Read the suggestion how to check reboot stability in https://progress.opensuse.org/projects/openqav3/wiki/Wiki#Best-practices-for-infrastructure-work• Try to reproduce the problem by rebooting openqaworker-arm-1 or openqaworker-arm-2 in a loop and check if the alert is triggered or pending for long enough so that the alert would trigger			
Related issues:			
Copied from openQA Infrastructure - action #92302: NFS mount var-lib-openqa-s...		Resolved	2021-06-11

History

#1 - 2021-05-23 07:54 - okurz

- Copied from action #92302: NFS mount var-lib-openqa-share.mount often fails after boot of some workers added

#2 - 2021-05-23 08:07 - okurz

On openqaworker-arm-3 in sudo journalctl -u os-autoinst-openvswitch.service:

```
May 22 01:43:24 openqaworker-arm-3 systemd[1]: Started os-autoinst openvswitch helper.
May 22 01:43:25 openqaworker-arm-3 os-autoinst-openvswitch[3238]: Waiting for IP on bridge 'br1', 60s left ...
May 22 01:43:48 openqaworker-arm-3 os-autoinst-openvswitch[3238]: Waiting for IP on bridge 'br1', 59s left ...
...
May 22 01:44:25 openqaworker-arm-3 os-autoinst-openvswitch[3238]: can't parse bridge local port IP at /usr/lib
/os-autoinst/os-autoinst-openvswitch line 46.
May 22 01:44:25 openqaworker-arm-3 os-autoinst-openvswitch[3238]: Waiting for IP on bridge 'br1', 1s left ...
```

so seems like <https://gitlab.suse.de/openqa/salt-states-openqa/-/commit/dbf659be290dad02d523502b7240dd4f0997e290> is not effective (anymore?) . https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/429#note_288016 states that a warning would be fixed but likely the original fix was not re-verified and got lost.

```
sudo systemctl show os-autoinst-openvswitch.service | grep OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT
```

shows Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=60 so not 300 as it should be.

Did sudo systemctl disable --now salt-minion telegraf openqa-worker-auto-restart@[1..100] (what was the better way again to stop all currently worker instances?)

We use an environment override file. Apparently it's not possible to override environment settings from an environment file <https://github.com/systemd/systemd/issues/9788> and an explicit environment file provided by option in <https://github.com/os-autoinst/os-autoinst/blob/master/systemd/os-autoinst-openvswitch.service.in#L15> is discouraged, see <https://github.com/systemd/systemd/issues/9788#issuecomment-420385947> but a systemd service override file as we use should work. I realized that the name of the directory is wrong, fix in

https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/495

Testing my changes on okurz-vm.qa with

```
export host=openqaworker-arm-3; for run in {01..30}; do for host in $host; do echo -n "run: $run, $host: ping .. " && timeout -k 5 600 sh -c "until ping -c30 $host >/dev/null; do :; done" && echo -n "ok, ssh .. " && timeout -k 5 600 sh -c "until nc -z -w 1 $host 22; do :; done" && echo -n "ok, uptime/reboot: " && ssh $host "uptime && sudo systemctl is-system-running && sudo reboot || sudo systemctl --failed" && sleep 120 || break; done || break; done
```

What I learned from this:

- Ask more diligently for at least *two* tests in all cases:
 - Is the new issue/feature working?
 - Does the old feature still work?
- Also be more diligent about testing reboot stability, see <https://progress.opensuse.org/projects/openqav3/wiki/#Best-practices-for-infrastructure-work>

#3 - 2021-05-23 14:53 - okurz

- Status changed from *Workable* to *Feedback*

- Assignee set to okurz

#4 - 2021-05-24 12:32 - okurz

- Status changed from *Feedback* to *Resolved*

```
> sudo salt -l error --no-color -C 'G@roles:worker' cmd.run 'sudo systemctl show os-autoinst-openvswitch.service | grep OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT'
openqaworker2.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker9.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker8.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
QA-Power8-4-kvm.qa.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
QA-Power8-5-kvm.qa.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
grenache-1.qa.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker5.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker6.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker3.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
powerqaworker-qam-1.qa.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker10.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker13.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
malbec.arch.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker-arm-1.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
openqaworker-arm-2.suse.de:
  Environment=OS_AUTOINST_USE_BRIDGE=br0 OS_AUTOINST_OPENVSWITCH_INIT_TIMEOUT=300
```

verified working fine on openqaworker-arm-3 over 40 reboots