

openQA Infrastructure - action #92302

NFS mount var-lib-openqa-share.mount often fails after boot of some workers

2021-05-07 09:52 - okurz

Status: Resolved	Start date:
Priority: High	Due date: 2021-06-11
Assignee: okurz	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: Ready	

Description

Observation

The nfs mount point systemd unit failed recently (and then turned to ok again) on one of our ARM workers. Likely the problem happens when an ARM machine is rebooted multiple times so that eventually we hit an alert window

Acceptance criteria

- **AC1:** No alert about failed systemd services related to NFS mount failing on ARM workers

Suggestions

- 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

Rollback

- ~~DONE: ssh openqaworker-arm-3 "sudo systemctl enable --now salt-minion"~~
- ~~DONE: ssh osd "salt key -y a openqaworker-arm-3.suse.de && sudo salt 'openqaworker-arm-3*' state.apply"~~
- ~~DONE: Unpause alerts for openqaworker-arm-3~~

Related issues:

Related to openQA Infrastructure - action #64941: after every reboot openqawo...	Resolved	2020-03-27	2021-06-11
Copied from openQA Infrastructure - action #89551: NFS mount fails after boot...	Resolved	2021-03-05	2021-03-31
Copied to openQA Infrastructure - action #92969: Failing service os-autoinst-...	Resolved	2021-05-23	

History

#1 - 2021-05-07 09:52 - okurz

- Copied from action #89551: NFS mount fails after boot (reproducible on some OSD workers) added

#2 - 2021-05-21 06:56 - Xiaojing_liu

openqaworker-arm-2 can't be reach using ping, so the osd deployment failed on 2021-05-21, but I didn't receive the alerting email.

#3 - 2021-05-23 07:42 - okurz

- Related to action #64941: after every reboot openqaworker7 is missing var-lib-openqa-share.mount , check dependencies of service with openqaworker1 added

#4 - 2021-05-23 07:49 - okurz

- Subject changed from NFS mount likely fails after boot of ARM workers to NFS mount var-lib-openqa-share.mount often fails after boot of some workers

- Priority changed from Normal to High

<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

os-autoinst-openvswitch tracked in [#92969](#)

#5 - 2021-05-23 07:54 - okurz

- Copied to action [#92969](#): Failing service os-autoinst-openvswitch after boot of some workers added

#6 - 2021-05-23 14:13 - okurz

- Status changed from Workable to In Progress

- Assignee set to okurz

[mkittler](#) can you comment about your suggestions in [#89551#note-8](#) . Have you tried them? Should we try again?

Trying on openqaworker-arm-3 with:

```
cat - > /etc/systemd/system/fix_mounts.service <<EOF
[Unit]
Description=Fix failed mounts by explicit mount command early (https://progress.opensuse.org/issues/92302)
Before=var-lib-openqa-share.mount
DefaultDependencies=no

[Service]
Type=oneshot
ExecStart=/usr/bin/mount -a

[Install]
WantedBy=multi-user.target
EOF
```

#7 - 2021-05-23 14:51 - okurz

- Status changed from In Progress to Feedback

https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/496

#8 - 2021-05-23 20:07 - okurz

- Assignee deleted (okurz)

fix_mounts.service fails after boot because it's also started to early:

```
May 23 19:47:44 openqaworker-arm-3 os-autoinst-openvswitch[2684]: Waiting for IP on bridge 'br1', 247s left ..
.
May 23 19:47:45 openqaworker-arm-3 mount[1147]: mount.nfs: Resource temporarily unavailable
May 23 19:47:45 openqaworker-arm-3 systemd[1]: fix_mounts.service: Main process exited, code=exited, status=64
/n/a
May 23 19:47:45 openqaworker-arm-3 systemd[1]: Failed to start Fix failed mounts by explicit mount command ear
ly (https://progress.opensuse.org/issues/92302).
May 23 19:47:45 openqaworker-arm-3 systemd[1]: fix_mounts.service: Unit entered failed state.
May 23 19:47:45 openqaworker-arm-3 systemd[1]: fix_mounts.service: Failed with result 'exit-code'.
May 23 19:47:45 openqaworker-arm-3 systemd[1]: Mounting /var/lib/openqa/share...
May 23 19:47:45 openqaworker-arm-3 systemd[1]: var-lib-openqa-share.mount: Mount process exited, code=exited s
tatus=32
May 23 19:47:45 openqaworker-arm-3 systemd[1]: Failed to mount /var/lib/openqa/share.
...
```

trying to introduce additional dependencies waiting for sockets.target.

#9 - 2021-05-23 20:07 - okurz

- Assignee set to okurz

#10 - 2021-05-24 12:36 - okurz

- Description updated

Tried increased timeout and retries in /etc/fstab

```
# extended NFS retry https://progress.opensuse.org/issues/89551
openqa.suse.de:/var/lib/openqa/share          /var/lib/openqa/share  nfs ro,x-systemd.mount-timeout=30m,retry=3
0 0 0
```

but this does not prevent the mount points to show up as failed for the first five minutes after boot. Also looks like this for the specific mount unit:

```
# journalctl -f -u var-lib-openqa-share.mount
May 24 14:23:49 openqaworker-arm-3 systemd[1]: Mounting /var/lib/openqa/share...
May 24 14:23:49 openqaworker-arm-3 systemd[1]: var-lib-openqa-share.mount: Mount process exited, code=exited s
tatus=32
May 24 14:23:49 openqaworker-arm-3 systemd[1]: Failed to mount /var/lib/openqa/share.
May 24 14:23:49 openqaworker-arm-3 systemd[1]: var-lib-openqa-share.mount: Unit entered failed state.
May 24 14:28:14 openqaworker-arm-3 systemd[1]: Mounting /var/lib/openqa/share...
May 24 14:28:16 openqaworker-arm-3 systemd[1]: Mounted /var/lib/openqa/share.
```

Trying again with fix_mounts.service on top:

```
# /etc/systemd/system/fix_mounts.service
[Unit]
Description=Fix failed mounts by explicit mount command early (https://progress.opensuse.org/issues/92302)
After=sockets.target
Wants=sockets.target
Before=var-lib-openqa-share.mount
DefaultDependencies=no

[Service]
Type=oneshot
ExecStart=/usr/bin/mount -a

[Install]
WantedBy=multi-user.target
```

this both combined seems to work good:

```
run: 01, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 14:37:26 up 0:15, 1 user, load averag
e: 1.09, 1.15, 0.93
running
Connection to openqaworker-arm-3 closed by remote host.
run: 02, openqaworker-arm-3: ping .. ok, ssh .. run: 03, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reb
oot: "System is booting up. See pam_nologin(8)"
Connection closed by 10.160.0.85 port 22
run: 04, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 14:48:56 up 0:05, 0 users, load avera
ge: 1.41, 1.41, 0.71
starting
run: 05, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 14:51:25 up 0:08, 0 users, load avera
ge: 0.80, 1.17, 0.72
running
Connection to openqaworker-arm-3 closed by remote host.
run: 06, openqaworker-arm-3: ping .. ok, ssh .. run: 07, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reb
oot: 15:03:12 up 0:05, 0 users, load average: 1.84, 1.47, 0.72
starting
run: 08, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:05:40 up 0:08, 0 users, load avera
ge: 0.31, 0.98, 0.65
running
Connection to openqaworker-arm-3 closed by remote host.
run: 09, openqaworker-arm-3: ping .. ok, ssh .. run: 10, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reb
oot: 15:16:48 up 0:05, 0 users, load average: 2.09, 1.47, 0.69
starting
run: 11, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:19:17 up 0:07, 0 users, load avera
ge: 1.22, 1.28, 0.73
running
Connection to openqaworker-arm-3 closed by remote host.
run: 12, openqaworker-arm-3: ping .. ok, ssh .. run: 13, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reb
oot: "System is booting up. See pam_nologin(8)"
Connection closed by 10.160.0.85 port 22
run: 14, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:30:47 up 0:05, 0 users, load avera
ge: 1.28, 1.33, 0.67
starting
run: 15, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:33:17 up 0:08, 0 users, load avera
ge: 0.68, 1.09, 0.68
running
Connection to openqaworker-arm-3 closed by remote host.
run: 16, openqaworker-arm-3: ping .. ok, ssh .. run: 17, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reb
```

```
oot: 15:44:22 up 0:05, 0 users, load average: 1.81, 1.37, 0.66
starting
run: 18, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:46:51 up 0:07, 0 users, load average: 1.04, 1.14, 0.67
running
Connection to openqaworker-arm-3 closed by remote host.
run: 19, openqaworker-arm-3: ping .. ok, ssh .. run: 20, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 15:57:58 up 0:05, 0 users, load average: 2.06, 1.40, 0.64
starting
run: 21, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:00:27 up 0:07, 1 user, load average: 1.27, 1.32, 0.73
running
```

so at least 14 successful reboots, the other runs either aborted or accessed the system when it was still in the process of bootup. Tried again with disabled `fix_mounts.service` and that is not enough:

```
echo "### disable fix_mounts.service again, only timeout and retry options in fstab"; 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 --no-legend" && sleep 120 || break; done || break; done
### disable fix_mounts.service again, only timeout and retry options in fstab
run: 01, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:21:40 up 0:15, 1 user, load average: 0.81, 0.84, 0.72
running
Connection to openqaworker-arm-3 closed by remote host.
run: 02, openqaworker-arm-3: ping .. ok, ssh .. run: 03, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:33:20 up 0:05, 0 users, load average: 1.58, 1.48, 0.76
starting
var-lib-openqa-share.mount loaded failed failed /var/lib/openqa/share
run: 04, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:35:49 up 0:08, 0 users, load average: 0.56, 1.13, 0.73
running
Connection to openqaworker-arm-3 closed by remote host.
run: 05, openqaworker-arm-3: ping .. ok, ssh .. run: 06, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:46:50 up 0:05, 0 users, load average: 2.04, 1.50, 0.71
starting
var-lib-openqa-share.mount loaded failed failed /var/lib/openqa/share
run: 07, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 16:49:20 up 0:07, 0 users, load average: 0.75, 1.15, 0.70
running
Connection to openqaworker-arm-3 closed by remote host.
run: 08, openqaworker-arm-3: ping .. ok, ssh .. run: 09, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 17:00:22 up 0:05, 0 users, load average: 1.72, 1.36, 0.65
starting
var-lib-openqa-share.mount loaded failed failed /var/lib/openqa/share
run: 10, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 17:02:50 up 0:07, 0 users, load average: 0.88, 1.18, 0.69
running
Connection to openqaworker-arm-3 closed by remote host.
run: 11, openqaworker-arm-3: ping .. ok, ssh .. run: 12, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 17:13:51 up 0:05, 0 users, load average: 1.75, 1.39, 0.67
starting
var-lib-openqa-share.mount loaded failed failed /var/lib/openqa/share
run: 13, openqaworker-arm-3: ping .. ok, ssh .. ok, uptime/reboot: 17:16:21 up 0:07, 0 users, load average: 0.86, 1.19, 0.71
running
Connection to openqaworker-arm-3 closed by remote host.
```

failing service `var-lib-openqa-share.mount`. The systemd documentation <https://www.freedesktop.org/software/systemd/man/systemd.mount.html#> explains "Network mount units automatically acquire `After=` dependencies on `remote-fs-pre.target`, `network.target` and `network-online.target`, and gain a `Before=` dependency on `remote-fs.target` unless `nofail` mount option is set. Towards the latter a `Wants=` unit is added as well." which I can confirm. But <https://www.freedesktop.org/wiki/Software/systemd/NetworkTarget/> states "By default all remote mounts defined in `/etc/fstab` pull this service in, in order to make sure the network is up before it is attempted to connect to a network share." which does not seem to work in our case. I suspect an actual bug in wicked. But as I assume no one is willing or able to fix that without us having a small and easy reproducer I am not even bothering reporting that elsewhere.

As alternatives I should try other options in `fstab`, e.g. `nofail,x-systemd.device-timeout=600,noauto,x-systemd.automount`

new experiment with `fstab` entry

```
# extended NFS retry https://progress.opensuse.org/issues/89551
openqa.suse.de:/var/lib/openqa/share /var/lib/openqa/share nfs ro,,noauto,nofail,retry=30,x-systemd.mount-timeout=30m,x-systemd.device-timeout=10m,x-systemd.automount 0 0
```

and no fix_mounts.service.

Running test with

```
echo -e "### disable fix_mounts.service again, fstab entry noauto,nofail,retry=30,x-systemd.mount-timeout=30m,x-systemd.device-timeout=10m,x-systemd.automount"; export host=openqaworker-arm-3; for run in {01..30}; do for host in $host; do echo -e -n "\nrn: $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 && test -d /var/lib/openqa/share/factory/ && sudo reboot || sudo systemctl --failed --no-legend" && echo -n " .. ok" && sleep 120 || break; done || break; done
```

20 reboots successful, https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/496 ready

#11 - 2021-05-24 21:52 - okurz

- *Description updated*

#12 - 2021-05-27 12:31 - okurz

- *Due date set to 2021-06-11*

checking after next weekend(s) if this prevents further alerts.

#13 - 2021-06-09 20:57 - okurz

- *Status changed from Feedback to Resolved*

<https://stats.openqa-monitor.qa.suse.de/d/KToPYLEWz/failed-systemd-services?orgId=1&from=now-30d&to=now> shows no reference to failed NFS mount units since long.