

openQA Infrastructure - action #89551

NFS mount fails after boot (reproducible on some OSD workers)

2021-03-05 12:17 - mkittler

Status:	Resolved	Start date:	2021-03-05
Priority:	High	Due date:	2021-03-31
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
problem			
<ul style="list-style-type: none">The NFS mount for /var/lib/openqa/share fails when booting. This has been reproduced on openqaworker13 (see #88900), openqaworker2 (see #89551#note-5) and possibly openqaworker-arm-2 (see #75016).Technically all workers might be affected because the problem is quite generic: The systemd service for the NFS mount does not wait until the ethernet connection is established.			
impact			
The mounting is actually re-attempted automatically after a few minutes so it shouldn't be a big deal. However, since the systemd unit for the mount stays failed for a few minutes false alerts are triggered which should be prevented.			
acceptance criteria			
<ul style="list-style-type: none">AC1: No false alerts are triggered if the NFS mountpoint fails just for a few minutes after booting.			
notes			
<ul style="list-style-type: none">The dependency to network-online.target is correct. However, network-online.target does not really ensure the ethernet connection is up in our setup. This is pretty clear from the logs mentioned in #88900#note-5.To configure network-online.target correctly further configuration would be required.<ul style="list-style-type: none">The official documentation only mentions how to do it for NetworkManager and systemd-networkd: https://www.freedesktop.org/wiki/Software/systemd/NetworkTarget/#cutthecraphowdoimakesurethatmyservicestartsafterthenetworkisreallyonlineLikely it doesn't make much sense relying on network-online.target anyways and just cope with the fact that the mountpoint might fail for a few minutes.			
Related issues:			
Related to openQA Infrastructure - action #88191: openqaworker2 boot ends in...		Resolved	2021-01-25
Related to openQA Infrastructure - action #68053: powerqaworker-qam-1 fails t...		Resolved	2020-06-14
Related to openQA Infrastructure - action #75016: [osd-admins][alert] Failed ...		Resolved	2020-10-21
Related to openQA Infrastructure - action #88900: openqaworker13 was unreachable		Resolved	2021-02-22
Copied to openQA Infrastructure - action #92302: NFS mount var-lib-openqa-sha...		Resolved	2021-06-11

History

#1 - 2021-03-05 12:31 - okurz

- Priority changed from Normal to High
- Target version set to Ready

As I assume this issue can cause confusing alerts and problems observable in other tickets as well I regard this as high priority and will reduce prio of other related tickets such as [#68053](#) and [#75016](#) and [#88191](#) accordingly

#2 - 2021-03-05 12:31 - okurz

- Related to action #88191: openqaworker2 boot ends in emergency shell added

#3 - 2021-03-05 12:32 - okurz

- Related to action #68053: powerqaworker-qam-1 fails to come up on reboot (repeatedly) added

#4 - 2021-03-05 12:32 - okurz

- Related to action #75016: [osd-admins][alert] Failed systemd services alert (workers): os-automount-openvswitch.service (and var-lib-openqa-share.mount) on openqaworker-arm-2 and others added

#5 - 2021-03-05 14:11 - mkittler

Also happening on openqaworker2:

```
[ OK ] Reached target Network is Online.
        Starting Reboot Manager...
        Mounting /var/lib/openqa/share...
        Starting Login and scanning of iSCSI devices...
        Starting Load kdump kernel and initrd...
[FAILED] Failed to mount /var/lib/openqa/share.
See 'systemctl status var-lib-openqa-share.mount' for details.
[DEPEND] Dependency failed for Remote File Systems.
[ OK ] Started OpenQA Worker Cache Service.
[ OK ] Started OpenQA Worker Cache Service Minion.
        Starting Permit User Sessions...
[ OK ] Started Login and scanning of iSCSI devices.
[ OK ] Started Permit User Sessions.
[ OK ] Started OpenSSH Daemon.
[ OK ] Started Reboot Manager.
[ OK ] Started Daemon to remotely execute Nagios plugins.
[ OK ] Started NTP Server Daemon.
[ OK ] Reached target System Time Synchronized.
[ OK ] Started Check if mainboard battery is Ok.
[ OK ] Started Backup of RPM database.
[ OK ] Started Discard unused blocks on a mounted filesystem.
[ OK ] Started Timeline of Snapper Snapshots.
[ OK ] Started daily update of the root trust anchor for DNSSEC.
[ OK ] Started Reminder for degraded MD arrays.
[ OK ] Started Nightly trigger of auto-update..
[ OK ] Started MD array scrubbing.
[ OK ] Started Command Scheduler.
[ OK ] Started Scrub btrfs filesystem, verify block checksums.
[ OK ] Started Balance block groups on a btrfs filesystem.
[ OK ] Started Daily rotation of log files.
[ OK ] Started Backup of /etc/sysconfig.
[ OK ] Started Discard unused blocks once a week.
[ OK ] Reached target Timers.
[ OK ] Started Load kdump kernel and initrd.
[ 77.476580] br0: port 1(eth0) entered blocking state
[ 77.482321] br0: port 1(eth0) entered disabled state
[ 77.488220] device eth0 entered promiscuous mode
[ 77.516255] device tap0 entered promiscuous mode
[ 77.565105] device tap1 entered promiscuous mode
[ 77.583359] device tap10 entered promiscuous mode
[ 77.606978] device tap11 entered promiscuous mode
[ 77.620830] device tap12 entered promiscuous mode
[ 77.645374] device tap128 entered promiscuous mode
[ 77.662705] device tap129 entered promiscuous mode
[ 77.686300] device tap13 entered promiscuous mode
[ 77.704775] device tap130 entered promiscuous mode
[ 77.725619] device tap131 entered promiscuous mode
[ 77.743597] device tap132 entered promiscuous mode
[ 77.765323] device tap133 entered promiscuous mode
[ 77.784885] device tap134 entered promiscuous mode
[ 77.817649] device tap135 entered promiscuous mode
[ 77.834431] device tap136 entered promiscuous mode
[ 77.855503] device tap137 entered promiscuous mode
[ 77.873705] device tap138 entered promiscuous mode
[ 77.900904] device tap139 entered promiscuous mode
[ 77.919075] device tap14 entered promiscuous mode
[ 77.942040] device tap140 entered promiscuous mode
[ 77.960235] device tap141 entered promiscuous mode
[ 77.982801] device tap142 entered promiscuous mode
[ 78.001112] device tap143 entered promiscuous mode
[ OK ] Started tgt admin.
[ 78.026556] device tap144 entered promiscuous mode
[ 78.039552] device tap145 entered promiscuous mode
[ 78.061487] device tap146 entered promiscuous mode
[ OK ] Started Update cron periods from /etc/sysconfig/btrfsmaintenance.
[ 78.079166] device tap147 entered promiscuous mode
```

```

[ 78.102144] device tap148 entered promiscuous mode
[ 78.120427] device tap149 entered promiscuous mode
[ 78.142676] device tap15 entered promiscuous mode
[ 78.160766] device tap150 entered promiscuous mode
[ 78.181627] device tap151 entered promiscuous mode
[ 78.200834] device tap152 entered promiscuous mode
[ 78.225611] device tap153 entered promiscuous mode
[ 78.241284] device tap154 entered promiscuous mode
[ 78.265979] device tap155 entered promiscuous mode
[ 78.285660] device tap16 entered promiscuous mode
[ 78.308208] device tap17 entered promiscuous mode
[ 78.327345] device tap18 entered promiscuous mode
[ 78.349424] device tap19 entered promiscuous mode
[ 78.369036] device tap2 entered promiscuous mode
[ 78.391967] device tap20 entered promiscuous mode
[ 78.411652] device tap21 entered promiscuous mode
[ 78.447049] device tap22 entered promiscuous mode
[ 78.466855] device tap23 entered promiscuous mode
[ 78.490362] device tap24 entered promiscuous mode
[ 78.510129] device tap25 entered promiscuous mode
[ 78.581411] device tap26 entered promiscuous mode
[ 78.603601] device tap27 entered promiscuous mode
[ 78.626346] device tap3 entered promiscuous mode
[ 78.646515] device tap4 entered promiscuous mode
[ 78.668644] device tap5 entered promiscuous mode
[ 78.687668] device tap6 entered promiscuous mode
[ 78.709956] device tap64 entered promiscuous mode
[ 78.729177] device tap65 entered promiscuous mode
[ 78.752471] device tap66 entered promiscuous mode
[ 78.771921] device tap67 entered promiscuous mode
[ 78.794848] device tap68 entered promiscuous mode
[ 78.814014] device tap69 entered promiscuous mode
[ 78.841224] device tap7 entered promiscuous mode
[ 78.861324] device tap70 entered promiscuous mode
[ 78.885475] device tap71 entered promiscuous mode
[ 78.906416] device tap72 entered promiscuous mode
[ 78.930531] device tap73 entered promiscuous mode
[ 78.950659] device tap74 entered promiscuous mode
[ 78.975211] device tap75 entered promiscuous mode
[ 78.997893] device tap76 entered promiscuous mode
[ 79.022463] device tap77 entered promiscuous mode
OK ] Started The Salt Minion.
[ 79.046314] device tap78 entered promiscuous mode
[ 79.070128] device tap79 entered promiscuous mode
[ 79.091225] device tap8 entered promiscuous mode
[ 79.116235] device tap80 entered promiscuous mode
[ 79.136992] device tap81 entered promiscuous mode
[ 79.162396] device tap82 entered promiscuous mode
[ 79.183811] device tap83 entered promiscuous mode
[ 79.210520] device tap84 entered promiscuous mode
[ 79.232698] device tap85 entered promiscuous mode
[ 79.269581] device tap86 entered promiscuous mode
[ 79.291124] device tap87 entered promiscuous mode
[ 79.316877] device tap88 entered promiscuous mode
[ 79.338280] device tap89 entered promiscuous mode
[ 79.363663] device tap9 entered promiscuous mode
[ 79.386391] device tap90 entered promiscuous mode
[ 79.412213] device tap91 entered promiscuous mode
[ 81.301252] igb 0000:01:00.0 eth0: igb: eth0 NIC Link is Up 1000 Mbps Full Duplex, Flow Control: RX
[ 81.310951] br0: port 1(eth0) entered blocking state
[ 81.316148] br0: port 1(eth0) entered forwarding state
[ 81.321625] IPv6: ADDRCONF(NETDEV_CHANGE): br0: link becomes ready
OK ] Stopped OpenQA Worker Cache Service Minion.
OK ] Started OpenQA Worker Cache Service Minion.
OK ] Started os-autoinst openvswitch helper.

```

Also happening on previous reboots:

```

martchus@openqaworker2:~> sudo journalctl -u var-lib-openqa-share.mount
-- Logs begin at Wed 2021-02-24 17:41:48 CET, end at Fri 2021-03-05 15:08:59 CET. --
Feb 28 03:31:30 openqaworker2 systemd[1]: Unmounting /var/lib/openqa/share...
Feb 28 03:31:30 openqaworker2 systemd[1]: Unmounted /var/lib/openqa/share.
-- Reboot --
Feb 28 03:36:08 openqaworker2 systemd[1]: Mounting /var/lib/openqa/share...
Feb 28 03:36:08 openqaworker2 systemd[1]: var-lib-openqa-share.mount: Mount process exited, code=exited status

```

```
=32
Feb 28 03:36:08 openqaworker2 systemd[1]: Failed to mount /var/lib/openqa/share.
Feb 28 03:36:08 openqaworker2 systemd[1]: var-lib-openqa-share.mount: Unit entered failed state.
Feb 28 03:40:33 openqaworker2 systemd[1]: Mounting /var/lib/openqa/share...
Feb 28 03:40:34 openqaworker2 systemd[1]: Mounted /var/lib/openqa/share.
Mär 05 14:57:26 openqaworker2 systemd[1]: Unmounting /var/lib/openqa/share...
Mär 05 14:57:26 openqaworker2 systemd[1]: Unmounted /var/lib/openqa/share.
-- Reboot --
Mär 05 15:03:29 openqaworker2 systemd[1]: Mounting /var/lib/openqa/share...
Mär 05 15:03:29 openqaworker2 systemd[1]: var-lib-openqa-share.mount: Mount process exited, code=exited status
=32
Mär 05 15:03:29 openqaworker2 systemd[1]: Failed to mount /var/lib/openqa/share.
Mär 05 15:03:29 openqaworker2 systemd[1]: var-lib-openqa-share.mount: Unit entered failed state.
Mär 05 15:07:54 openqaworker2 systemd[1]: Mounting /var/lib/openqa/share...
Mär 05 15:07:55 openqaworker2 systemd[1]: Mounted /var/lib/openqa/share.
```

This also confirms that it automatically tries again after a few minutes. I suppose we can rely on that.

#6 - 2021-03-05 14:12 - mkittler

- Description updated

- Assignee set to mkittler

#7 - 2021-03-07 21:14 - okurz

- Related to action #88900: openqaworker13 was unreachable added

#8 - 2021-03-10 15:44 - mkittler

Maybe some mount options can help (https://man.archlinux.org/man/nfs.5#MOUNT_OPTIONS,
<https://man.archlinux.org/man/systemd.mount.5#FSTAB>).

If we would add `x-systemd.mount-timeout=10m,retry=10` it would actually retry for 10 minutes and only then be considered failed. If a try times out after 3 minutes as the documentation suggests this would mean at least 3 tries. (We need `x-systemd.mount-timeout=10m` to prevent systemd from stopping/failing the mount earlier if `DefaultTimeoutStartSec=` is set to a lower timeout in the systemd config.)

#9 - 2021-03-12 10:58 - okurz

- Status changed from New to In Progress

#10 - 2021-03-12 12:53 - mkittler

Alternatively we could also just rely on https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/463 which triggers the alert only after 15 minutes.

#11 - 2021-03-17 14:37 - okurz

With the above MR merged, how would you see the current situation?

#12 - 2021-03-17 15:04 - okurz

- Due date set to 2021-03-31

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

#13 - 2021-03-29 12:13 - mkittler

- Status changed from In Progress to Feedback

With the above MR merged, how would you see the current situation?

I actually haven't seen any *related* systemd service alerts anymore.

#14 - 2021-03-29 12:25 - okurz

- Status changed from Feedback to Resolved

I can confirm that now `sudo salt '*' cmd.run 'systemctl --no-legend --failed'` does not show the problem and AFAIR no related alert was raised recently. As discussed with mkittler setting to "Resolved".

#15 - 2021-05-07 09:52 - okurz

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