

openQA Infrastructure - action #78010

unreliable reboots on openqaworker3, likely due do openqa_nvme_format (was: [alert] PROBLEM Host Alert: openqaworker3.suse.de is DOWN)

2020-11-16 12:48 - okurz

Status: Resolved	Start date: 2020-11-16
Priority: High	Due date: 2021-04-21
Assignee: okurz	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: Ready	

Description

Observation

alert by email:
From: Monitoring User nagios@suse.de resent from: okurz@suse.com
To: okurz@suse.com
Date: 16/11/2020 10.01
Spam Status: Spamassassin
Notification: PROBLEM
Host: openqaworker3.suse.de
State: DOWN
Date/Time: Mon Nov 16 09:01:00 UTC 2020
Info: CRITICAL - 10.160.0.243: Host unreachable @ 10.160.0.44. rta nan, lost 100%

See Online: <https://thruk.suse.de/thruk/cgi-bin/extinfo.cgi?type=1host=openqaworker3.suse.de>

Acceptance criteria

- **AC1:** openqaworker3 is "reboot-safe", e.g. at least 10 reboots in a row end up in a successfully booted system

Related issues:

Related to openQA Infrastructure - action #68050: openqaworker3 fails to come...	Resolved	2020-06-14	2020-07-07
Related to openQA Infrastructure - action #71098: openqaworker3 down but no a...	Resolved	2020-09-08	
Related to openQA Infrastructure - action #88385: openqaworker3 host up alert...	Rejected	2021-02-01	
Related to openQA Infrastructure - action #88191: openqaworker2 boot ends in...	Resolved	2021-01-25	

History

#1 - 2020-11-16 12:49 - okurz

- Due date set to 2020-11-18

- Status changed from Workable to Feedback

I checked right now with ping openqaworker3 and could ping the system. and ssh openqaworker uptime showed that the system is up for 0:41 , so since about 12:07 UTC.

everything seems to be in order.

#2 - 2020-11-16 12:59 - nicksinger

Guess that was me. A simple "chassis power cycle" did the trick.

#3 - 2020-11-16 15:20 - okurz

Interesting. But do you know what *caused* the problem then?

#4 - 2020-11-19 16:10 - okurz

- Assignee changed from okurz to nicksinger

[nicksinger](#) Interesting. But do you know what *caused* the problem then?

#5 - 2020-11-24 10:24 - nicksinger

- Assignee changed from nicksinger to okurz

no. SOL was stuck once again and I just triggered a reboot

#6 - 2020-11-25 21:08 - okurz

- Status changed from Feedback to Resolved

- Assignee changed from okurz to nicksinger

hm, ok. I checked again, everything seems to be in order. I don't know what else we can do. Ok, thanks for fixing it.

#7 - 2020-11-29 14:32 - okurz

- Related to action #68050: openqaworker3 fails to come up on reboot, openqa_nvme_format.service failed added

#8 - 2020-11-29 14:33 - okurz

- Related to action #71098: openqaworker3 down but no alert was raised added

#9 - 2020-11-29 14:42 - okurz

- Subject changed from [alert] PROBLEM Host Alert: openqaworker3.suse.de is DOWN to unreliable reboots on openqaworker3, likely due do openqa_nvme_format (was: [alert] PROBLEM Host Alert: openqaworker3.suse.de is DOWN)

- Description updated

- Due date deleted (2020-11-18)

- Status changed from Resolved to Workable

- Priority changed from Urgent to High

ok. I know what we can do. This failed again after the weekly automatic reboot whenever there are kernel updates. openqaworker3 was stuck in emergency mode and also a power reset this time ended up in a similar situation. It could be that it would work to just reboot once more but as we had problems repeatedly and still have we should test at least 5-10 reboots in a row.

In sol activate I found:

```
[FAILED] Failed to start Setup NVMe before mounting it.
See 'systemctl status openqa_nvme_format.service' for details.
[DEPEND] Dependency failed for /var/lib/openqa.
[DEPEND] Dependency failed for Local File Systems.
[DEPEND] Dependency failed for /var/lib/openqa/share.
[DEPEND] Dependency failed for Remote File Systems.
[DEPEND] Dependency failed for openQA Worker #7.
[DEPEND] Dependency failed for openQA Worker #12.
[DEPEND] Dependency failed for openQA Worker #8.
[DEPEND] Dependency failed for openQA Worker #5.
[DEPEND] Dependency failed for openQA Worker #2.
[DEPEND] Dependency failed for openQA Worker #3.
[DEPEND] Dependency failed for openQA Worker #9.
[DEPEND] Dependency failed for openQA Worker #1.
[DEPEND] Dependency failed for openQA Worker #13.
[DEPEND] Dependency failed for openQA Worker #11.
[DEPEND] Dependency failed for openQA Worker #4.
[DEPEND] Dependency failed for openQA Worker #6.
[DEPEND] Dependency failed for openQA Worker #10.
```

Seems in [#68050](#) I could not make that properly.

#10 - 2020-11-29 15:00 - coolo

but also ssh service failed

#11 - 2020-11-30 13:59 - nicksinger

I've rejected the salt-key for now on OSD to prevent automatic startup of workers. What I found while booting is that the systems hangs for quite some time with the last message being printed: [1034.526144] kexec_file: kernel signature verification failed (-129). - which seems to fit with [okurz](#) suggestion that this is caused by kernel updates.

#12 - 2020-12-03 20:55 - okurz

nicksinger wrote:

[...] seems to fit with [okurz](#) suggestion that this is caused by kernel updates.

but what I meant is only that some package upgrades, for example kernel updates, trigger a reboot which is all according to plan. Have you seen [#78010#note-9](#) regarding the systemd services? I suspect that this is simply again or still a problem of systemd service dependencies and fully in our control.

#13 - 2020-12-04 08:29 - okurz

please see [#77011#note-18](#) for changes that fvogt has applied. We should crosscheck the changes he did, commit to salt and ensure this is applicable for all machines and then apply the same to all, o3 and osd.

#14 - 2020-12-04 09:25 - nicksinger

So I diffed with what fvogt did:

```
diff --git a/openqa/nvme_store/openqa_nvme_create.service b/openqa/nvme_store/openqa_nvme_create.service
new file mode 100644
index 0000000..9e9b57b
--- /dev/null
+++ b/openqa/nvme_store/openqa_nvme_create.service
@@ -0,0 +1,20 @@
+[Unit]
+Description=Create array on NVMe if necessary
+# Let's hope this is close enough to "all nvmes present"
+Requires=dev-nvme0n1.device
+After=dev-nvme0n1.device
+DefaultDependencies=no
+
+[Service]
+Type=oneshot
+
+# It's not really possible to wait for that to happen, so do it here
+ExecStart=/bin/sh -c "if ! mdadm --detail --scan | grep -qi openqa; then mdadm --assemble --scan || true; fi"
+# Create striped storage for openQA from all NVMe devices when / resides on
+# another device or from a potential third NVMe partition when there is only a
+# single NVMe device for the complete storage
+# For some reason mdadm --create doesn't send an udev event, so do it manually.
+
+ExecStart=/bin/sh -c -e 'if lsblk -n | grep -q "raid"; then exit 0; fi; if lsblk -n | grep -v nvme | grep -q "/"
+/$"; then mdadm --create /dev/md/openqa --level=0 --force --raid-devices=$(ls /dev/nvme?n1 | wc -l) --run /dev
+/nvme?n1; else mdadm --create /dev/md/openqa --level=0 --force --raid-devices=1 --run /dev/nvme0n1p3; fi; udev
+adm trigger -c add /dev/md/openqa'
+
+[Install]
+WantedBy=multi-user.target
diff --git a/openqa/nvme_store/openqa_nvme_format.service b/openqa/nvme_store/openqa_nvme_format.service
index f6f1c16..8e9d97c 100644
--- a/openqa/nvme_store/openqa_nvme_format.service
+++ b/openqa/nvme_store/openqa_nvme_format.service
@@ -1,19 +1,12 @@
 [Unit]
-Description=Setup NVMe before mounting it
+Description=Create Ext2 FS on /dev/md/openqa
+Before=var-lib-openqa.mount
+Requires=dev-md-openqa.device
+After=dev-md-openqa.device
+DefaultDependencies=no

[Service]
Type=oneshot
-
-# Create striped storage for openQA from all NVMe devices when / resides on
-# another device or from a potential third NVMe partition when there is only a
-# single NVMe device for the complete storage
-
+ExecStart=/bin/sh -c 'lsblk -n | grep -q "raid" || lsblk -n | grep -v nvme | grep "/" && (mdadm --stop /dev/m
+d/openqa >/dev/null 2>&1; mdadm --create /dev/md/openqa --level=0 --force --raid-devices=$(ls /dev/nvme?n1 | w
+c -l) --run /dev/nvme?n1) || mdadm --create /dev/md/openqa --level=0 --force --raid-devices=1 --run /dev/nvme0
+n1p3'
-# Ensure device is correctly initialized but also spend a little time before
-# trying to create a filesystem to prevent a "busy" error
-ExecStart=/bin/sh -c 'grep nvme /proc/mdstat'
-ExecStart=/bin/sh -c 'mdadm --detail --scan | grep openqa'
+ExecStart=/sbin/mkfs.ext2 -F /dev/md/openqa
```

[Install]

I'm not convinced yet that this will resolve our problem. I'd rather like to look into why our array-devices are busy in the first place. I *suspect* that something assembles our array before we format and therefore assume we have a dependency problem.

From a first, quick look I think it could be related to mdmonitor.service. From <https://linux.die.net/man/8/mdadm>:

```
[...] all arrays listed in the configuration file will be monitored. Further, if --scan is given, then any other md devices that appear in /proc/mdstat will also be monitored.
```

Which would explain why the device is busy. I will experiment with some dependencies for our service to run before mdmonitor

#15 - 2020-12-11 12:42 - nicksinger

- Status changed from Workable to Blocked

Hm, I had the bright idea of removing the mdraid module from the initramfs. This now causes the machine to hang in the dracut recovery shell (so even earlier than we had before). I wanted to rebuild the initramfs but failed to boot any recovery media. I've opened [RT-ADM #182375] AutoReply: Machine openqaworker3.suse.de boots wrong PXE image now to address the fact that PXE on that machine seems to boot some infra machine image. And of course I forgot to add osd-admins@suse.de into CC so I will keep you updated here...

#16 - 2020-12-15 04:00 - nicksinger

- Status changed from Blocked to In Progress

machine can boot from PXE again. Apparently we had specific PXE images for each openqaworker back in the days, interesting :) I was able to boot a TW installer in which I can hopefully recover the initramfs again.

#17 - 2020-12-15 04:10 - nicksinger

well, seems like this custom image was (blindly? Or was I pressing Enter where no console showed up?) reinstalling the worker. At least this is how it looks to me:

```
0:openqaworker3:~ # lsblk
NAME          MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
loop0         7:0    0  89.1M  1 loop  /parts/mp_0000
loop1         7:1    0  12.7M  1 loop  /parts/mp_0001
loop2         7:2    0  58.3M  1 loop  /mounts/mp_0000
loop3         7:3    0  72.6M  1 loop  /mounts/mp_0001
loop4         7:4    0   4.1M  1 loop  /mounts/mp_0002
loop5         7:5    0   1.8M  1 loop  /mounts/mp_0003
sda           8:0    0  931.5G  0 disk
├─sda1        8:1    0   9.8G  0 part
│  └─md4      9:4    0   9.8G  0 raid1 /mnt
├─sda2        8:2    0  995.6M  0 part
└─sda3        8:3    0  920.8G  0 part
   └─md0      9:0    0  920.8G  0 raid1
sdb           8:16   0  931.5G  0 disk
├─sdb1        8:17   0   9.8G  0 part
│  └─md4      9:4    0   9.8G  0 raid1 /mnt
├─sdb2        8:18   0  995.6M  0 part
└─sdb3        8:19   0  920.8G  0 part
   └─md0      9:0    0  920.8G  0 raid1
nvme0n1       259:0   0  372.6G  0 disk
└─nvme0n1p1  259:1   0  372.6G  0 part
nvme1n1       259:2   0  372.6G  0 disk
0:openqaworker3:~ # ls -lah /mnt/
total 132K
drwxr-xr-x 27 root root 4.0K Dec 11 12:15 .
drwxr-xr-x 23 root root  820 Dec 14 13:34 ..
drwxr-xr-x  2 root root 4.0K Apr 13  2018 bin
drwxr-xr-x  3 root root 4.0K Dec 11 12:15 boot
-rw-r--r--  1 root root  893 Apr 13  2018 bootincluded_archives.filelist
-rw-r--r--  1 root root  816 May 24  2013 build-custom
drwxr-xr-x  3 root root 4.0K Apr 13  2018 config
drwxr-xr-x  3 root root 4.0K Apr 13  2018 dev
drwxr-xr-x 90 root root 4.0K Dec 14 11:34 etc
drwxr-xr-x  2 root root 4.0K Jun 27  2017 home
drwxr-xr-x  2 root root 4.0K Dec 11 12:14 kiwi-hooks
drwxr-xr-x  2 root root 4.0K Apr 13  2018 kvm
drwxr-xr-x  2 root root 4.0K Apr 13  2018 kvm_lock_sync
drwxr-xr-x 10 root root 4.0K Apr 13  2018 lib
drwxr-xr-x  7 root root 4.0K Apr 13  2018 lib64
```

```

drwx----- 2 root root 16K Apr 13 2018 lost+found
drwxr-xr-x 2 root root 4.0K Jun 27 2017 mnt
drwxr-xr-x 2 root root 4.0K Jun 27 2017 opt
drwxr-xr-x 2 root root 4.0K Apr 13 2018 proc
drwx----- 4 root root 4.0K Apr 13 2018 root
drwxr-xr-x 16 root root 4.0K Apr 13 2018 run
drwxr-xr-x 2 root root 12K Dec 11 12:14 sbin
drwxr-xr-x 2 root root 4.0K Jun 27 2017 selinux
drwxr-xr-x 5 root root 4.0K Apr 13 2018 srv
drwxr-xr-x 3 root root 4.0K May 24 2013 studio
dr-xr-xr-x 2 root root 4.0K Jun 27 2017 sys
drwxrwxrwt 9 root root 4.0K Dec 14 11:34 tmp
drwxr-xr-x 13 root root 4.0K Apr 13 2018 usr
drwxr-xr-x 11 root root 4.0K Dec 11 12:15 var

```

Guess it's time for a reinstall...

#18 - 2020-12-15 07:49 - nicksinger

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

the re-installation went quite smooth. After the initial install was done I re-added it into salt and applied a highstate, after that a reboot, adjusted /etc/salt/grains to include the worker and nvme_store role, ran another highstate, another reboot and with that got the NVME's as raid for /var/lib/openqa and all MM interfaces/bridges up and running.

A first glimpse at <https://stats.openqa-monitor.qa.suse.de/d/4KkGdvZk/osd-status-overview?orgId=1> as well as <https://stats.openqa-monitor.qa.suse.de/alerting/list> looks good (no alerts). Also a test on public cloud looks good till now: <https://openqa.suse.de/tests/5173202#>

#19 - 2021-01-19 11:28 - nicksinger

- Status changed from *Feedback* to *Resolved*

#20 - 2021-01-25 02:58 - Xiaojing_liu

On 2021-1-25 morning, openqaworker3 failed to start. I log in to this machine use ipmitool and press Control-D, and it showed:

```
[SOL Session operational. Use ~? for help]
```

```
Login incorrect
```

```
Give root password for maintenance
```

```
(or press Control-D to continue):
```

```
Starting default target
```

```
[86888.689121] openvswitch: Open vSwitch switching datapath
```

```
[86890.320765] No iBFT detected.
```

```
Welcome to openSUSE Leap 15.2 - Kernel 5.3.18-lp152.60-default (ttyS1).
```

```
openqaworker3 login: [86891.801204] bridge: filtering via arp/ip/ip6tables is no longer available by default.
```

```
Update your scripts to load br_netfilter if you need this.
```

```
[86891.825522] device ovs-system entered promiscuous mode
```

```
[86891.831409] Timeout policy base is empty
```

```
[86891.835347] Failed to associated timeout policy `ovs_test_tp'
```

```
[86892.109539] device br1 entered promiscuous mode
```

```
[86892.235381] tun: Universal TUN/TAP device driver, 1.6
```

```
[86892.798181] gre: GRE over IPv4 demultiplexor driver
```

```
[86892.855322] ip_gre: GRE over IPv4 tunneling driver
```

```
[86892.861814] device gre_sys entered promiscuous mode
```

```
[86895.245432] bpfILTER: Loaded bpfILTER_umh pid 2680
```

```
[86895.245646] Started bpfILTER
```

```
[86903.635371] BTRFS info (device sda3): balance: start -dusage=0
```

```
[86903.641304] BTRFS info (device sda3): balance: ended with status: 0
```

```
[86906.584005] BTRFS info (device sda3): balance: start -dusage=5
```

```
[86906.590673] BTRFS info (device sda3): balance: ended with status: 0
```

```
...
```

#21 - 2021-01-26 11:40 - okurz

- Status changed from *Resolved* to *Feedback*

Well, I can confirm that one can login into the machine with ssh but still openqa_nvme_format reports as failed, journal-output:

```
Jan 25 03:42:21 openqaworker3 sh[1090]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
```

```
Jan 25 03:42:21 openqaworker3 systemd[1]: Starting Setup NVMe before mounting it...
```

```
Jan 25 03:42:21 openqaworker3 sh[1090]: mdadm: cannot open /dev/nvme0n1p3: No such file or directory
Jan 25 03:42:21 openqaworker3 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status=2/INVALIDARGUMENT
Jan 25 03:42:21 openqaworker3 systemd[1]: Failed to start Setup NVMe before mounting it.
Jan 25 03:42:21 openqaworker3 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Jan 25 03:42:21 openqaworker3 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
```

so not really "reliable reboots" which we would expect here. [nicksinger](#) please update from your side or set back to "Workable" and unassign.

#22 - 2021-02-01 11:51 - cdywan

- Related to action #88385: openqaworker3 host up alert is flaky added

#23 - 2021-02-15 11:19 - cdywan

- Observed a "host up alert" today, *ssh* would time out, hit ^D via *ipmi*
- *ssh* claims it's not the expected host anymore and my key is denied... apparently I hit another corner case, which I worked around by explicitly specifying the key... `□□`
- `systemctl list-units --failed w/` no obvious way to fix i.e. restarting the service won't work

#24 - 2021-02-16 09:47 - cdywan

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

#25 - 2021-02-22 02:58 - Xiaojing_liu

On 2021-02-22, openqaworker3 failed to reboots (was: [Alerting] openqaworker3: host up alert). just like <https://progress.opensuse.org/issues/78010#note-20> said.

openqa_nvme_format.service failed on openqaworker3 after it rebooted success. (was: [Alerting] Failed systemd services alert (except openqa.suse.de)),

Checking in: <https://stats.openqa-monitor.qa.suse.de/d/KToPYLEWz/failed-systemd-services?orgId=1>

The result in openqaworker3 is:

```
xiaojing@openqaworker3:~> sudo systemctl --failed
UNIT                                LOAD    ACTIVE SUB    DESCRIPTION
● openqa_nvme_format.service loaded failed failed Setup NVMe before mounting it
```

LOAD = Reflects whether the unit definition was properly loaded.
ACTIVE = The high-level unit activation state, i.e. generalization of SUB.
SUB = The low-level unit activation state, values depend on unit type.

1 loaded units listed. Pass `--all` to see loaded but inactive units, too.
To show all installed unit files use `'systemctl list-unit-files'`.

#26 - 2021-03-15 05:58 - okurz

openqaworker3 not ready today, blocking osd-deployment. Manually continued boot process with `^D`

#27 - 2021-04-06 09:28 - okurz

- Status changed from Feedback to In Progress

- Assignee changed from [nicksinger](#) to [okurz](#)

<https://gitlab.suse.de/openqa/salt-states-openqa/-/commit/7c3192df54e9d9e368ec1fda866357486f706e76> introduced a problem for openqaworker8+9 which both repeatedly fail to reboot as now the RAID creation arguments would be setup before stopping so that the `lsblk` output is misinterpreted when `/dev/md/openqa` is still present. Conducted debugging session [okurz+mkittler](#). On openqaworker8 manually masked relevant openQA services, *telegraf* and *salt* and paused monitoring. Manually adapted scripts on openqaworker8, [mkittler](#) will create a MR. Currently running experiment with multiple reboot attempts in a screen session belonging to root on osd. After this can test on other machines, e.g. openqaworker2 and re-enable openqaworker8.

#28 - 2021-04-06 09:40 - mkittler

Mentioned MR: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/474

#29 - 2021-04-06 11:56 - mkittler

Looks like it succeeded 10 reboots. I've just started another round of reboots.

#30 - 2021-04-07 04:16 - openqa_review

- Due date set to 2021-04-21

#31 - 2021-04-08 13:12 - okurz

openqaworker8 is fully online again. Unpaused the corresponding host-up alert in grafana. Next candidate for testing could be again openqaworker3 as for this ticket. On openqaworker3 did `systemctl mask --now telegraf salt-minion openqa-worker-cacheservice && systemctl mask --now openqa-worker-auto-restart@{1..13}`

checked stability in a loop with

```
for run in {01..30}; do for host in openqaworker3; 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 reboot" && sleep 120 || break; done || break; done
```

and all 30 reboots were successful.

In the journal we can check if the retry was triggered:

```
# journalctl --since=today -u openqa_nvme_format.service | grep 'Trying RAID0 creation again'
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: Trying RAID0 creation again after timeout (attempt 2 of 10)
Apr 08 11:43:03 openqaworker3 openqa-establish-nvme-setup[920]: Trying RAID0 creation again after timeout (attempt 2 of 10)
Apr 08 13:33:37 openqaworker3 openqa-establish-nvme-setup[904]: Trying RAID0 creation again after timeout (attempt 2 of 10)
Apr 08 13:39:54 openqaworker3 openqa-establish-nvme-setup[903]: Trying RAID0 creation again after timeout (attempt 2 of 10)
Apr 08 14:16:00 openqaworker3 openqa-establish-nvme-setup[901]: Trying RAID0 creation again after timeout (attempt 2 of 10)
```

with full logs from the first entry:

```
Apr 08 11:03:12 openqaworker3 systemd[1]: Starting Setup NVME before mounting it...
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: Current mount points (printed for debugging purposes):
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: devtmpfs on /dev type devtmpfs (rw,nosuid,size=131889732k,nr_inodes=32972433,mode=755)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: tmpfs on /run type tmpfs (rw,nosuid,nodev,mode=755)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/unified type cgroup2 (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,xattr,name=sys>)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: none on /sys/fs/bpf type bpf (rw,nosuid,nodev,noexec,relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,perf_event)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,memory)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,devices)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/rdma type cgroup (rw,nosuid,nodev,noexec,relatime,rdma)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,freezer)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,cpuset)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/pids type cgroup (rw,
```

```

nosuid,nodev,noexec,relatime,pids)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/net_cls,net_prio type
cgroup (rw,nosuid,nodev,noexec,relatime,net_c<
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/blkio type cgroup (rw
,nosuid,nodev,noexec,relatime,blkio)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: cgroup on /sys/fs/cgroup/cpu,cpuacct type cgro
up (rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on / type btrfs (rw,relatime,space_c
ache,subvolid=268,subvol=@/.snapshots/1/snaps>
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: systemd-1 on /proc/sys/fs/binfmt_misc type aut
ofs (rw,relatime,fd=34,pgrp=1,timeout=0,minpro>
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: debugfs on /sys/kernel/debug type debugfs (rw,
relatime)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: mqueue on /dev/mqueue type mqueue (rw,relatime
)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: hugetlbfs on /dev/hugepages type hugetlbfs (rw
,relatime,pagesize=2M)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /.snapshots type btrfs (rw,relati
me,space_cache,subvolid=267,subvol=@/.snapsho>
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /boot/grub2/x86_64-efi type btrfs
(rw,relatime,space_cache,subvolid=265,subvol=>
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /boot/grub2/i386-pc type btrfs (r
w,relatime,space_cache,subvolid=266,subvol=@/>
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /tmp type btrfs (rw,relatime,spac
e_cache,subvolid=260,subvol=@/tmp)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /root type btrfs (rw,relatime,spa
ce_cache,subvolid=262,subvol=@/root)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /usr/local type btrfs (rw,relatim
e,space_cache,subvolid=259,subvol=@/usr/local)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /opt type btrfs (rw,relatime,spac
e_cache,subvolid=263,subvol=@/opt)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /home type btrfs (rw,relatime,spa
ce_cache,subvolid=264,subvol=@/home)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /srv type btrfs (rw,relatime,spac
e_cache,subvolid=261,subvol=@/srv)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: /dev/sda3 on /var type btrfs (rw,relatime,spac
e_cache,subvolid=258,subvol=@/var)
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: Present block devices (printed for debugging p
urposes):
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: NAME          MAJ:MIN RM   SIZE RO TYPE  MOUNTPOINT
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: sda             8:0    0 931.5G  0 disk
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sda1          8:1    0   9.8G  0 part
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: | └md126       9:126  0   9.8G  0 raid1
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sda2          8:2    0 995.6M  0 part
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: | └md125       9:125  0 995.5M  0 raid1
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sda3          8:3    0 920.8G  0 part /
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: sdb             8:16   0 931.5G  0 disk
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sdb1          8:17   0   9.8G  0 part
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: | └md126       9:126  0   9.8G  0 raid1
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sdb2          8:18   0 995.6M  0 part
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: | └md125       9:125  0 995.5M  0 raid1
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └sdb3          8:19   0 920.8G  0 part
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: nvme1n1       259:0   0 372.6G  0 disk
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └md127         9:127  0  745G  0 raid0
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: nvme0n1       259:1   0 372.6G  0 disk
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: └md127         9:127  0  745G  0 raid0
Apr 08 11:03:12 openqaworker3 openqa-establish-nvme-setup[901]: Creating RAID0 "/dev/md/openqa" on: /dev/nvme0
n1 /dev/nvme1n1
Apr 08 11:03:15 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: cannot open /dev/nvme0n1: Device or res
ource busy
Apr 08 11:03:15 openqaworker3 openqa-establish-nvme-setup[901]: Waiting 10 seconds before trying again after f
ailing due to busy device.
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: Trying RAID0 creation again after timeout (att
empt 2 of 10)
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: Stopping current RAID "/dev/md/openqa"
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: stopped /dev/md/openqa
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: /dev/nvme0n1 appears to be part of a ra
id array:
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]:          level=raid0 devices=2 ctime=Thu Apr 8
10:57:02 2021
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: partition table exists on /dev/nvme0n1
but will be lost or
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]:          meaningless after creating array
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: /dev/nvme1n1 appears to be part of a ra

```


id array:

```
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: level=raid0 devices=2 ctime=Thu Apr 8
10:57:02 2021
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: partition table exists on /dev/nvme1n1
but will be lost or
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: meaningless after creating array
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: Defaulting to version 1.2 metadata
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mdadm: array /dev/md/openqa started.
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: Status for RAID0 "/dev/md/openqa"
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: ARRAY /dev/md/openqa metadata=1.2 name=openqaw
orker3:openqa UUID=4984c770:c8a94fd6:488f2ccc:>
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: Creating ext2 filesystem on RAID0 "/dev/md/ope
nqa"
Apr 08 11:03:23 openqaworker3 openqa-establish-nvme-setup[901]: mke2fs 1.43.8 (1-Jan-2018)
Apr 08 11:03:25 openqaworker3 openqa-establish-nvme-setup[901]: [182B blob data]
Apr 08 11:03:25 openqaworker3 openqa-establish-nvme-setup[901]: Creating filesystem with 195289600 4k blocks a
nd 48824320 inodes
```

so this works but I wonder if we can find out what blocks the device. This is something I could check later again, maybe on the same machine. But as I enabled the services again let's give them a chance to finish :)

#32 - 2021-04-08 14:57 - okurz

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

actually after thinking about it I think this is good enough. AC1 is fulfilled and if the retry helps that should be good enough, not needing further investigation.