

openQA Infrastructure - action #88191

openqaworker2 boot ends in emergency shell

2021-01-25 13:50 - dzedro

Status: Resolved	Start date: 2021-01-25
Priority: Normal	Due date:
Assignee: mkittler	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: Ready	

Description

Observation

openqaworker2 was down since 24.1. 2:30 AM
I started the worker ignoring the reason why it ended in maintenance shell.
Looks like ext4/ext2 mount issue.
This should be fixed, next planed reboot will have same result.

Acceptance criteria

- **AC1:** openqaworker2 is up and running again
- **AC2:** openqaworker2 is stable over multiple reboots

Suggestions

- Trigger reboot explicitly multiple times and check that the machine comes up properly.

Related issues:

Related to openQA Infrastructure - action #78010: unreliable reboots on openq...	Resolved	2020-11-16	2021-04-21
Related to openQA Infrastructure - action #89551: NFS mount fails after boot ...	Resolved	2021-03-05	2021-03-31

History

#1 - 2021-01-25 20:25 - cdywan

- Status changed from New to Rejected

I think this was probably fixed in the meantime. I can't see any issues now

#2 - 2021-01-25 22:03 - dzedro

- Status changed from Rejected to New

Fixed by whom ? It is not fixed ...

#3 - 2021-01-26 08:57 - okurz

- Project changed from openQA Project to openQA Infrastructure

- Description updated

- Status changed from New to Workable

- Target version set to Ready

good in the case the issue is not present currently anymore but before being confident on this not appearing again I suggest to check reboot stability 2-3 times.

EDIT: I could login over ssh but systemctl --failed shows degraded and openqa_nvme_format.service as failed

#4 - 2021-02-16 09:47 - cdywan

- Related to action #78010: unreliable reboots on openqaworker3, likely due do openqa_nvme_format (was: [alert] PROBLEM Host Alert: openqaworker3.suse.de is DOWN) added

#5 - 2021-02-16 09:50 - cdywan

[#78010](#) seems to look very similar at this point, getting stuck in emergency and failing on nvme setup

[nicksinger](#) Maybe you could take a look at this

#6 - 2021-03-05 12:31 - okurz

- Related to action [#89551](#): NFS mount fails after boot (reproducible on some OSD workers) added

#7 - 2021-03-05 12:31 - okurz

- Priority changed from High to Normal

I suggest we work on [#89551](#) first which could be related to some problems on bootstrap

#8 - 2021-03-05 13:39 - mkittler

- Assignee set to mkittler

#9 - 2021-03-05 14:33 - mkittler

I could reboot one time successfully (except for running into [#89551](#)).

However, trying again I can reproduce the issue:

```
[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.
Starting Prepare NVMe after mounting it...
Starting Restore /run/initramfs on shutdown...
...
Welcome to openSUSE Leap 15.2!

[ OK ] Stopped Switch Root.
[ OK ] Stopped Journal Service.
Starting Journal Service...
Mounting POSIX Message Queue File System...
[ OK ] Listening on Process Core Dump Socket.
[ OK ] Listening on LVM2 poll daemon socket.
Starting Load Kernel Modules...
[ OK ] Listening on udev Control Socket.
[ OK ] Listening on Device-mapper event daemon FIFOs.
Starting Setup NVMe before mounting it...
[ OK ] Created slice system-getty.slice.
Starting Create list of required st...ce nodes for the current kernel...
Mounting Huge Pages File System...
[ OK ] Created slice system-serial\x2dgetty.slice.
[ OK ] Started Dispatch Password Requests to Console Directory Watch.
[ OK ] Reached target Local Encrypted Volumes.
[ OK ] Stopped File System Check on Root Device.
Starting Remount Root and Kernel File Systems...
Mounting Kernel Debug File System...
[ 16.080966] BTRFS info (device md1): disk space caching is enabled
Starting Monitoring of LVM2 mirrors...ng dmeventd or progress polling...
[ OK ] Listening on /dev/initctl Compatibility Named Pipe.
[ OK 16.121189] md127: detected capacity change from 0 to 799906201600
0m] Reached target Remote File Systems (Pre).
[ OK ] Set up automount Arbitrary Executab...rmats File System Automount Point.
[ OK ] Stopped target Switch Root.
[ OK ] Stopped target Initrd File Systems.
[ OK ] Stopped target Initrd Root File System.
[ OK ] Listening on udev Kernel Socket.
[ OK ] Created slice User and Session Slice.
[ OK ] Reached target Slices.
[ OK ] Mounted Kernel Debug File System.
[ OK ] Mounted POSIX Message Queue File System.
[ OK ] Mounted Huge Pages File System.
[ OK ] Started Journal Service.
[ OK ] Started Load Kernel Modules.
[ OK ] Started Create list of required sta...vice nodes for the current kernel.
[ OK ] Started Remount Root and Kernel File Systems.
```

```

[ OK ] Stopped Entropy Daemon based on the HAVEGE algorithm.
[ OK ] Started Entropy Daemon based on the HAVEGE algorithm.
        Starting udev Coldplug all Devices...
        Starting Load/Save Random Seed...
        Starting Create Static Device Nodes in /dev...
        Starting Apply Kernel Variables...
[ OK ] Started Load/Save Random Seed.
[ OK ] Started Apply Kernel Variables.
[ OK ] Started udev Coldplug all Devices.
        Starting udev Wait for Complete Device Initialization...
[ OK ] Started Create Static Device Nodes in /dev.
        Starting udev Kernel Device Manager...
[ OK ] Started udev Kernel Device Manager.
[ 17.196918] input: Power Button as /devices/LNXSYSTM:00/LNXPWRBN:00/input/input0
[ 17.206110] ACPI: Power Button [PWRB]
[ 17.210289] input: Power Button as /devices/LNXSYSTM:00/LNXPWRBN:00/input/input1
[ 17.218467] ACPI: Power Button [PWRB]
[ 17.247908] dca service started, version 1.12.1
[ 17.310929] igb: Intel(R) Gigabit Ethernet Network Driver - version 5.6.0-k
[ 17.318228] igb: Copyright (c) 2007-2014 Intel Corporation.
[ 17.393049] igb 0000:01:00.0: added PHC on eth0
[ 17.397917] igb 0000:01:00.0: Intel(R) Gigabit Ethernet Network Connection
[ 17.405146] igb 0000:01:00.0: eth0: (PCIe:5.0Gb/s:Width x4) 2c:60:0c:73:02:ac
[ 17.420405] igb 0000:01:00.0: eth0: PBA No: 106300-000
[ 17.425874] igb 0000:01:00.0: Using MSI-X interrupts. 8 rx queue(s), 8 tx queue(s)
[ 17.487081] igb 0000:01:00.1: added PHC on eth1
[ 17.491949] igb 0000:01:00.1: Intel(R) Gigabit Ethernet Network Connection
[ 17.499163] igb 0000:01:00.1: eth1: (PCIe:5.0Gb/s:Width x4) 2c:60:0c:73:02:ad
[ 17.506713] igb 0000:01:00.1: eth1: PBA No: 106300-000
[ 17.512184] igb 0000:01:00.1: Using MSI-X interrupts. 8 rx queue(s), 8 tx queue(s)
[ 18.444581] mpt3sas_cm0: port enable: SUCCESS
[ 18.449770] scsi 10:0:0:0: Enclosure QUANTA S2B 2U 2.5 0700 PQ: 0 ANSI: 5
[ 18.458554] scsi 10:0:0:0: set ignore_delay_remove for handle(0x000a)
[ 18.465325] scsi 10:0:0:0: SES: handle(0x000a), sas_addr(0x500605b0000272bd), phy(36), device_name(0x500605
b0000272bd)
[ 18.476618] scsi 10:0:0:0: enclosure logical id (0x500605b0000272bf), slot(0)
[ 18.484437] scsi 10:0:0:0: enclosure level(0x0000), connector name( )
[ 18.491563] scsi 10:0:0:0: qdepth(1), tagged(0), scsi_level(6), cmd_que(0)
[ 18.499035] end_device-10:0:0: add: handle(0x000a), sas_addr(0x500605b0000272bd)
[ 18.507419] scsi 10:0:0:0: Attached scsi generic sg2 type 13
[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.
        Starting Prepare NVMe after mounting it...
        Starting Restore /run/initramfs on shutdown...
[ OK ] Stopped Dispatch Password Requests to Console Directory Watch.
[ OK ] Reached target Login Prompts.
[ OK ] Reached target Network (Pre).
[ OK ] Reached target User and Group Name Lookups.
[ OK ] Reached target Timers.
[ OK ] Reached target System Time Synchronized.
[ OK ] Reached target Network.
[ OK ] Reached target Network is Online.
[ OK ] Reached target Host and Network Name Lookups.
[ OK ] Reached target Sockets.
[ OK ] Reached target Paths.
[ OK ] Started Emergency Shell.
[ OK ] Reached target Emergency Mode.
[ OK ] Started Restore /run/initramfs on shutdown.
[ OK ] Started Prepare NVMe after mounting it.
[ 22.322090] IPMI message handler: version 39.2
[ 22.336720] ipmi device interface
[ 22.398112] ipmi_si: IPMI System Interface driver
[ 22.403173] ipmi_si dmi-ipmi-si.0: ipmi_platform: probing via SMBIOS
[ 22.409912] ipmi_platform: ipmi_si: SMBIOS: io 0xca2 regsize 1 spacing 1 irq 0
[ 22.417847] ipmi_si: Adding SMBIOS-specified kcs state machine
[ 22.424146] ipmi_si IPI0001:00: ipmi_platform: probing via ACPI
[ 22.430547] ipmi_si IPI0001:00: ipmi_platform: [io 0xca2] regsize 1 spacing 1 irq 0
[ 22.439091] ipmi_si dmi-ipmi-si.0: Removing SMBIOS-specified kcs state machine in favor of ACPI
[ 22.448468] ipmi_si: Adding ACPI-specified kcs state machine
[ 22.454574] ipmi_si: Trying ACPI-specified kcs state machine at i/o address 0xca2, slave address 0x20, irq
0

```

```

[ 22.812597] ipmi_si IPI0001:00: IPMI message handler: Found new BMC (man_id: 0x001c4c, prod_id: 0x3242, dev
_id: 0x20)
[ 22.892804] ipmi_si IPI0001:00: IPMI kcs interface initialized
[ 22.942993] ipmi_ssif: IPMI SSIF Interface driver
[ 23.673070] i801_smbus 0000:00:1f.3: SMBus using PCI interrupt
[ 23.680309] mei_me 0000:00:16.0: Device doesn't have valid ME Interface
[ 23.793433] input: PC Speaker as /devices/platform/pcspkr/input/input2
[ 23.878182] cryptd: max_cpu_qlen set to 1000
[ 23.882997] ixgbe: Intel(R) 10 Gigabit PCI Express Network Driver - version 5.1.0-k
[ 23.891340] ixgbe: Copyright (c) 1999-2016 Intel Corporation.
[ 24.274254] [drm] Using P2A bridge for configuration
[ 24.279556] [drm] AST 2400 detected
[ 24.283399] [drm] Analog VGA only
[ 24.287080] [drm] dram MCLK=408 Mhz type=1 bus_width=16 size=00800000
[ 24.293925] [TTM] Zone kernel: Available graphics memory: 131899552 KiB
[ 24.300993] [TTM] Zone dma32: Available graphics memory: 2097152 KiB
[ 24.307868] [TTM] Initializing pool allocator
[ 24.312576] [TTM] Initializing DMA pool allocator
[ 24.322190] fbcon: astdrmfb (fb0) is primary device
[ 24.366968] Console: switching to colour frame buffer device 128x48
[ 24.395783] ast 0000:07:00.0: fb0: astdrmfb frame buffer device
[ 24.432597] [drm] Initialized ast 0.1.0 20120228 for 0000:07:00.0 on minor 0
[ 24.441206] AVX2 version of gcm_enc/dec engaged.
[ 24.445849] AES CTR mode by8 optimization enabled
[ OK ] Started Monitoring of LVM2 mirrors,...sing dmeventd or progress polling.
[ OK ] Reached target Local File Systems (Pre).
Mounting /var/lib/machines...
Mounting /boot/grub2/x86_64-efi...
Mounting /var/lib/mariadb...
Mounting /var/spool...
[ 24.651073] kvm: unknown parameter 'nested' ignored
Mounting /tmp...
Mounting /var/lib/pgsql...
Mounting /var/lib/named...
Mounting /usr/local...
Mounting /var/tmp...
Mounting /var/opt...
Mounting /var/lib/mailman...
Mounting /home...
Mounting /var/log...
Mounting /var/crash...
Mounting /var/lib/libvirt/images...
Mounting /.snapshots...
Mounting /boot/grub2/i386-pc...
Mounting /srv...
Mounting /opt...
Mounting /var/lib/mysql...
[ OK ] Mounted /var/tmp.
[ OK ] Mounted /var/spool.
[ OK ] Mounted /var/opt.
[ OK ] Mounted /var/log.
[ OK 24.925229] EDAC MC0: Giving out device to module sb_edac controller Haswell SrcID#1_Ha#0: DEV 0000:f
f:12.0 (INTERRUPT)
0m] Mounted /var[ 24.936992] EDAC MC1: Giving out device to module sb_edac controller Haswell SrcID#0_Ha#0:
DEV 0000:7f:12.0 (INTERRUPT)
[ 24.949049] EDAC sbridge: Ver: 1.1.2
/lib/pgsql.
[ OK ] Mounted /var/lib/named.
[ OK ] Mounted /var/lib/mysql.
[ OK ] Mounted /var/lib/mariadb.
[ OK ] Mounted /var/lib/mailman.
[ OK ] Mounted /var/lib/machines.
[ OK ] Mounted /var/lib/libvirt/images.
[ OK ] Mounted /var/crash.
[ OK ] Mounted /usr/local.
[ OK ] Mounted /tmp.
[ OK ] Mounted /srv.
[ 25.089092] ixgbe 0000:04:00.0: Multiqueue Enabled: Rx Queue count = 32, Tx Queue count = 32 XDP Queue coun
t = 0
[ OK 25.099581] ixgbe 0000:04:00.0: 32.000 Gb/s available PCIe bandwidth (5 GT/s x8 link)
0m] Mounted /opt[ 25.108661] ixgbe 0000:04:00.0: MAC: 2, PHY: 1, PBA No: FFFFFFFF-0FF
.
[ 25.116148] ixgbe 0000:04:00.0: 2c:60:0c:21:8c:ef
[ OK ] Mounted /home.
[ 25.124255] ixgbe 0000:04:00.0: Intel(R) 10 Gigabit Network Connection

```

```

[ 25.130948] libphy: ixgbe-mdio: probed
[ OK ] Mounted /boot/grub2/x86_64-efi.
[ OK ] Mounted /boot/grub2/i386-pc.
[ OK ] Mounted /.snapshots.
Starting Flush Journal to Persistent Storage...
[ 25.226534] systemd-journald[931]: Received request to flush runtime journal from PID 1
Starting Setup NVMe before mounting it...
[ 25.619601] md127: detected capacity change from 799906201600 to 0
[ 25.625808] md: md127 stopped.
[ 25.744771] nvme0n1: pl
[ 25.750664] md127: detected capacity change from 0 to 799906201600
[ 25.768888] iTCO_vendor_support: vendor-support=0
[ 25.796644] nvme1n1:
[ 25.829040] iTCO_wdt: Intel TCO WatchDog Timer Driver v1.11
[ 25.835344] iTCO_wdt: unable to reset NO_REBOOT flag, device disabled by hardware/BIOS
[ 25.871992] ses 10:0:0:0: Attached Enclosure device
[ 25.953336] intel_rapl_common: Found RAPL domain package
[ 25.959325] intel_rapl_common: Found RAPL domain dram
[ 25.965027] intel_rapl_common: DRAM domain energy unit 15300pj
[ 25.988500] intel_rapl_common: Found RAPL domain package
[ 25.994409] intel_rapl_common: Found RAPL domain dram
[ 25.999992] intel_rapl_common: DRAM domain energy unit 15300pj
Starting MD array monitor...
[ OK ] Started MD array monitor.
You are in emergency mode. After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to boot
into default mode.
[ 26.321180] ixgbe 0000:04:00.1: Multiqueue Enabled: Rx Queue count = 32, Tx Queue count = 32 XDP Queue coun
t = 0
[ 26.332254] ixgbe 0000:04:00.1: 32.000 Gb/s available PCIe bandwidth (5 GT/s x8 link)
[ 26.340730] ixgbe 0000:04:00.1: MAC: 2, PHY: 1, PBA No: FFFFFFF-OFF
[ 26.347449] ixgbe 0000:04:00.1: 2c:60:0c:21:8c:f0
[ 26.355824] ixgbe 0000:04:00.1: Intel(R) 10 Gigabit Network Connection
[ 26.362912] libphy: ixgbe-mdio: probed
Give root password for maintenance
(or press Control-D to continue): [ 26.396595] Adding 2103228k swap on /dev/md0. Priority:-2 extents:1 acro
ss:2103228k FS
[ 48.035522] EXT4-fs (md127): mounting ext2 file system using the ext4 subsystem
[ 48.048618] EXT4-fs (md127): mounted filesystem without journal. Opts: (null)

```

#10 - 2021-03-05 14:51 - mkittler

- Status changed from Workable to In Progress

When continuing in the emergency shell it continues to boot:

```

martchus@openqaworker2:~> sudo journalctl --boot -u openqa_nvme_format.service
-- Logs begin at Wed 2021-02-24 17:41:48 CET, end at Fri 2021-03-05 15:46:46 CET. --
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 05 15:17:44 openqaworker2 sh[935]: level=raid0 devices=2 ctime=Fri Mar 5 15:02:29 2021
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or
Mär 05 15:17:44 openqaworker2 sh[935]: meaningless after creating array
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: /dev/nvme1n1 appears to be part of a raid array:
Mär 05 15:17:44 openqaworker2 sh[935]: level=raid0 devices=2 ctime=Fri Mar 5 15:02:29 2021
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: partition table exists on /dev/nvme1n1 but will be lost or
Mär 05 15:17:44 openqaworker2 sh[935]: meaningless after creating array
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: Defaulting to version 1.2 metadata
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: array /dev/md/openqa started.
Mär 05 15:17:49 openqaworker2 sh[935]: mdadm: timeout waiting for /dev/md/openqa
Mär 05 15:17:49 openqaworker2 sh[1042]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 05 15:17:49 openqaworker2 sh[1043]: ARRAY /dev/md127 metadata=1.2 name=openqaworker2:openqa UUID=a4b70ab9:
45e0d31b:53543112:fed7b743
Mär 05 15:17:49 openqaworker2 mkfs.ext2[1046]: mke2fs 1.43.8 (1-Jan-2018)
Mär 05 15:17:49 openqaworker2 mkfs.ext2[1046]: The file /dev/md/openqa does not exist and no size was specifie
d.
Mär 05 15:17:49 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=1/FAILURE
Mär 05 15:17:49 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 05 15:17:49 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 05 15:17:49 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 05 15:17:54 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 05 15:17:57 openqaworker2 sh[1168]: level=raid0 devices=2 ctime=Fri Mar 5 15:17:44 2021
Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or
Mär 05 15:17:57 openqaworker2 sh[1168]: meaningless after creating array

```

```

Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: /dev/nvme1n1 appears to be part of a raid array:
Mär 05 15:17:57 openqaworker2 sh[1168]: level=raid0 devices=2 ctime=Fri Mar 5 15:17:44 2021
Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: partition table exists on /dev/nvme1n1 but will be lost or
Mär 05 15:17:57 openqaworker2 sh[1168]: meaningless after creating array
Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: Defaulting to version 1.2 metadata
Mär 05 15:17:57 openqaworker2 sh[1168]: mdadm: array /dev/md/openqa started.
Mär 05 15:17:57 openqaworker2 sh[1230]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: mke2fs 1.43.8 (1-Jan-2018)
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: [220B blob data]
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: Creating filesystem with 195289600 4k blocks and 48824320 inode
s
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: Filesystem UUID: ffa08f78-bc30-40ba-ba60-9e666f1721a4
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: Superblock backups stored on blocks:
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1
605632, 2654208,
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: 4096000, 7962624, 11239424, 20480000, 23887872, 7166361
6, 78675968,
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: 102400000
Mär 05 15:17:57 openqaworker2 mkfs.ext2[1240]: [65B blob data]
Mär 05 15:17:57 openqaworker2 sh[1231]: ARRAY /dev/md/openqa metadata=1.2 name=openqaworker2:openqa UUID=27b02
fbf:440e1dda:eb216b83:2b9e8017
Mär 05 15:18:01 openqaworker2 mkfs.ext2[1240]: [134B blob data]
Mär 05 15:18:16 openqaworker2 mkfs.ext2[1240]: [99B blob data]
Mär 05 15:18:16 openqaworker2 systemd[1]: Started Setup NVMe before mounting it.
Mär 05 15:36:52 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 05 15:36:52 openqaworker2 sh[1296]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
Mär 05 15:36:52 openqaworker2 sh[1296]: mdadm: cannot open /dev/nvme0n1p3: No such file or directory
Mär 05 15:36:52 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=2/INVALIDARGUMENT
Mär 05 15:36:52 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 05 15:36:52 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 05 15:36:52 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 05 15:41:53 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 05 15:41:53 openqaworker2 sh[5496]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
Mär 05 15:41:53 openqaworker2 sh[5496]: mdadm: cannot open /dev/nvme0n1p3: No such file or directory
Mär 05 15:41:53 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=2/INVALIDARGUMENT
Mär 05 15:41:53 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 05 15:41:53 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 05 15:41:53 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.

```

Apparently it fails waiting for /dev/md/openqa. The 2nd attempt fails as well but that's only because at this point /dev/nvme0n1 is already used.

#11 - 2021-03-08 12:01 - mkittler

Just as reference, the service is defined here:

https://gitlab.suse.de/openqa/salt-states-openqa/-/blob/master/openqa/nvme_store/openqa_nvme_format.service

Crucial part of the log with comments:

```

# ExecStart=... mdadm --create ...
Mär 05 15:17:44 openqaworker2 sh[935]: mdadm: array /dev/md/openqa started.
Mär 05 15:17:49 openqaworker2 sh[935]: mdadm: timeout waiting for /dev/md/openqa # seems to contradict the pre
vious statement but likely the source of the problem
# ExecStart=/bin/sh -c 'grep nvme /proc/mdstat'
Mär 05 15:17:49 openqaworker2 sh[1042]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
# ExecStart=/bin/sh -c 'mdadm --detail --scan | grep openqa'
Mär 05 15:17:49 openqaworker2 sh[1043]: ARRAY /dev/md127 metadata=1.2 name=openqaworker2:openqa UUID=a4b70ab9:
45e0d31b:53543112:fed7b743
# ExecStart=/sbin/mkfs.ext2 -F /dev/md/openqa
Mär 05 15:17:49 openqaworker2 mkfs.ext2[1046]: mke2fs 1.43.8 (1-Jan-2018)
Mär 05 15:17:49 openqaworker2 mkfs.ext2[1046]: The file /dev/md/openqa does not exist and no size was specifie
d.
Mär 05 15:17:49 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=1/FAILURE

```

A good run looks like this (no timeout):

```

martchus@openqaworker2:~> sudo journalctl --boot -u openqa_nvme_format.service
-- Logs begin at Thu 2021-02-25 21:31:37 CET, end at Mon 2021-03-08 12:41:14 CET. --
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 05 15:58:40 openqaworker2 sh[943]: level=raid0 devices=2 ctime=Fri Mar 5 15:17:54 2021
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or
Mär 05 15:58:40 openqaworker2 sh[943]: meaningless after creating array
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: /dev/nvme1n1 appears to be part of a raid array:

```

```

Mär 05 15:58:40 openqaworker2 sh[943]: level=raid0 devices=2 ctime=Fri Mar 5 15:17:54 2021
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: partition table exists on /dev/nvme1n1 but will be lost or
Mär 05 15:58:40 openqaworker2 sh[943]: meaningless after creating array
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: Defaulting to version 1.2 metadata
Mär 05 15:58:40 openqaworker2 sh[943]: mdadm: array /dev/md/openqa started.
Mär 05 15:58:41 openqaworker2 sh[1051]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 05 15:58:42 openqaworker2 sh[1052]: ARRAY /dev/md/openqa metadata=1.2 name=openqaworker2:openqa UUID=84e79
eae:c458c2f6:b24308b1:afb22670
Mär 05 15:58:42 openqaworker2 mkfs.ext2[1055]: mke2fs 1.43.8 (1-Jan-2018)
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: [182B blob data]
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: Creating filesystem with 195289600 4k blocks and 48824320 inode
s
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: Filesystem UUID: f91bd8d9-e67c-414f-829e-86c7a8a4f33b
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: Superblock backups stored on blocks:
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1
605632, 2654208,
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: 4096000, 7962624, 11239424, 20480000, 23887872, 7166361
6, 78675968,
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: 102400000
Mär 05 15:58:44 openqaworker2 mkfs.ext2[1055]: [65B blob data]
Mär 05 15:58:49 openqaworker2 mkfs.ext2[1055]: [152B blob data]
Mär 05 15:59:04 openqaworker2 mkfs.ext2[1055]: [99B blob data]
Mär 05 15:59:04 openqaworker2 systemd[1]: Started Setup NVMe before mounting it.

```

I'm not sure what causes the timeout. Note that the mdadm --create command seems to succeed despite the timeout (exists with 0 return code). Otherwise systemd would not continue executing the next command. So I suspect that if we wanted to retry we couldn't just rely on mdadm's return code. This also means that (at least for this particular problem) the existing ||-fallback within the unit file's code is likely not effective.

Maybe it would already help to add a retry on system level. It is not possible to simply restart the unit later (see log in previous comment) but I'd assume if we'd add Restart=on-failure (and e.g. RestartSec=1, StartLimitBurst=10, StartLimitIntervalSec=1min). It would not be a nice workaround but better than being stuck in the rescue shell.

#12 - 2021-03-08 12:20 - cdywan

mkittler wrote:

Just as reference, the service is defined here:

https://gitlab.suse.de/openqa/salt-states-openqa/-/blob/master/openqa/nvme_store/openqa_nvme_format.service

In case you're not aware, see [419](#) for an attempt at reworking this to be more generic - but I left the MR aside for the moment due to lack of CI coverage so I wasn't very confident in making such changes.

#13 - 2021-03-08 14:53 - mkittler

I'm testing the SR created for [#80482](#) (also see [#80482#note-14](#)) and will try to extend it by the retry parameters mentioned in the previous comment.

#14 - 2021-03-08 17:14 - mkittler

It looks like restarting with the help of systemd isn't as easy as expected. I've tried:

```

[Unit]
...
StartLimitBurst=10
StartLimitIntervalSec=1min

[Service]
Type=oneshot
RestartSec=1

```

However, it fails with:

```

Mär 08 17:36:42 openqaworker2 systemd[1]: openqa_nvme_format.service: Service has Restart= setting other than
no, which isn't allowed for Type=oneshot services. Refusing.

```

Apparently systemd on the worker is too old:

```

martchus@openqaworker2:~> systemd --version
systemd 234

```

We needed at least version 244: <https://github.com/systemd/systemd/pull/13754#issuecomment-763146803>

Unfortunately the worker is currently busted by this. I can not simply revert the changed to the service file because the filesystem is read-only. Maybe it boots again RW if I disable the service or I'll mount the file system RW manually.

#15 - 2021-03-09 06:55 - mkittler

Looks like Leap 15.3 will have systemd 246: <https://build.opensuse.org/package/show/openSUSE:Leap:15.3/systemd> - Should we just wait for that?

#16 - 2021-03-09 07:20 - okurz

You mean wait for that to resolve the problem? No, as I don't trust openSUSE Leap 15.3 will be stable enough for very soon given that even no one cared to add proper migration tests before calling it Beta. Also, 15.3 is a bigger shift from 15.2 than just 15.1 to 15.2 with the complete change of how packages are provided from SLE.

#17 - 2021-03-10 10:58 - mkittler

About the strangeness of "array started" followed by a timeout: The relevant function is in <https://git.kernel.org/pub/scm/utils/mdadm/mdadm.git/tree/Create.c> and it is quite big with lots of gotos. I suppose the "started" is just referring to telling the kernel via some async ioctl commands to create the array. The timeout then occurs in wait_for() in <https://git.kernel.org/pub/scm/utils/mdadm/mdadm.git/tree/util.c#n1132>. That function is basically just polling with stat(). The attempts/delay for the polling is hard-coded so we can not change that. Considering that wait_for() is merely logging an error in the end it is also no surprise that the problem isn't reflected in mdadms return code. From the log timestamps we can tell that wait_for() only times out after 5 minutes. Likely there's no use in adding further delays from our side without completely retrying the creation.

#18 - 2021-03-10 11:03 - okurz

Care to open an upstream bug report about that problem then?

mkittler wrote:

[...] Likely there's no use in adding further delays from our side without completely retrying the creation.

Yes, looks like it. So, do you plan to do exactly that: Retry the creation?

#19 - 2021-03-10 11:14 - mkittler

Our version of mdadm also seems a bit old:

```
martchus@openqa-worker2:~> sudo mdadm --version
mdadm - v4.1 - 2018-10-01
```

However, that's actually the latest release.

I'm wondering whether updating the Kernel to >= 5.4 might be a problem in the future because there are recent Git commits which might be relevant then: <https://git.kernel.org/pub/scm/utils/mdadm/mdadm.git/commit/Create.c?id=329dfc28debb58ffe7bd1967cea00fc583139aca> (Btw, we're currently using Linux 5.3.)

#20 - 2021-03-10 11:17 - mkittler

I can try to open an upstream bug.

Yes, looks like it. So, do you plan to do exactly that: Retry the creation?

As a first step I'd move the commands from the systemd file into a shell script (because the retry on systemd-level doesn't work in our version). Then I'd try to implement retrying the creation.

#21 - 2021-03-10 12:58 - okurz

mkittler wrote:

Our version of mdadm also seems a bit old:

```
martchus@openqa-worker2:~> sudo mdadm --version
mdadm - v4.1 - 2018-10-01
```

However, that's actually the latest release.

I'm wondering whether updating the Kernel to >= 5.4 might be a problem in the future because there are recent Git commits which might be relevant then: <https://git.kernel.org/pub/scm/utils/mdadm/mdadm.git/commit/Create.c?id=329dfc28debb58ffe7bd1967cea00fc583139aca> (Btw, we're currently using Linux 5.3.)

If you can reproduce the error in a clean setup, e.g. have a proper openQA test or a manual VM environment then consider reporting this as bug against Leap 15.2 (and hence implicitly SLE15-SP2) as well. At best with a crosscheck against Tumbleweed which has both a more recent mdadm as well as Linux Kernel. Don't forget, SUSE is also building the OS we are using ;)

#22 - 2021-03-10 16:33 - mkittler

#23 - 2021-03-11 17:15 - mkittler

When testing https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/462 it is mostly as before. I can still reproduce the issue and sometimes it works. (Considering the SR request so far just moves the code this is expected.) However, I've also seen one occurrence of a different error:

```

martchus@openqaworker2:~> sudo journalctl --boot -u openqa_nvme_format.service
-- Logs begin at Tue 2021-03-02 09:49:36 CET, end at Thu 2021-03-11 15:49:42 CET. --
Mär 11 15:40:59 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 11 15:40:59 openqaworker2 openqa-establish-nvme-setup[1141]: Creating RAID0 "/dev/md/openqa" on: /dev/
nvme0n1 /dev/nvme1n1
Mär 11 15:40:59 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, st
atus=2/INVALIDARGUMENT
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1141]: mdadm: cannot open /dev/nvme0n1: Device o
r resource busy
Mär 11 15:40:59 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: Stopping current RAID "/dev/md/openqa"
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: stopped /dev/md/openqa
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: Creating RAID0 "/dev/md/openqa" on: /dev/
nvme0n1 /dev/nvme1n1
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: /dev/nvme0n1 appears to be part of
a raid array:
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]:          level=raid0 devices=2 ctime=Thu Ma
r 11 15:27:03 2021
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: partition table exists on /dev/nvm
e0n1 but will be lost or
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]:          meaningless after creating array
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: /dev/nvme1n1 appears to be part of
a raid array:
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]:          level=raid0 devices=2 ctime=Thu Ma
r 11 15:27:03 2021
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: partition table exists on /dev/nvm
e1n1 but will be lost or
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]:          meaningless after creating array
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: Defaulting to version 1.2 metadata
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mdadm: array /dev/md/openqa started.
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: Status for RAID0 "/dev/md/openqa"
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: md127 : active raid0 nvme1n1[1] nvme0n1[0
]
Mär 11 15:40:59 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: ARRAY /dev/md/openqa metadata=1.2 name=op
enqaworker2:openqa UUID=c524d3cc:01c420e8:8539dddf:cadc46ba
Mär 11 15:40:59 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: Creating ext2 filesystem on RAID0 "/dev/m
d/openqa"
Mär 11 15:41:02 openqaworker2 openqa-establish-nvme-setup[1170]: mke2fs 1.43.8 (1-Jan-2018)
Mär 11 15:41:00 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]: [182B blob data]
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]: Creating filesystem with 195289600 4k blo
cks and 48824320 inodes
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]: Filesystem UUID: cb6da583-f812-42a4-a37a-
74e8a71b8e89
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]: Superblock backups stored on blocks:
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]:          32768, 98304, 163840, 229376, 294
912, 819200, 884736, 1605632, 2654208,
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]:          4096000, 7962624, 11239424, 20480
000, 23887872, 71663616, 78675968,
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]:          102400000
Mär 11 15:41:03 openqaworker2 openqa-establish-nvme-setup[1170]: [65B blob data]
Mär 11 15:41:07 openqaworker2 openqa-establish-nvme-setup[1170]: [134B blob data]
Mär 11 15:41:22 openqaworker2 openqa-establish-nvme-setup[1170]: [99B blob data]
Mär 11 15:41:22 openqaworker2 systemd[1]: Started Setup NVMe before mounting it.
Mär 11 15:43:19 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 11 15:43:19 openqaworker2 openqa-establish-nvme-setup[1306]: Stopping current RAID "/dev/md/openqa"
Mär 11 15:43:19 openqaworker2 openqa-establish-nvme-setup[1306]: mdadm: Cannot get exclusive access to /de
v/md/openqa:Perhaps a running process, mounted filesystem or active volume group?
Mär 11 15:43:20 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, st
atus=1/FAILURE
Mär 11 15:43:20 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 11 15:43:20 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 11 15:43:20 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 11 15:48:19 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 11 15:48:19 openqaworker2 openqa-establish-nvme-setup[5013]: Stopping current RAID "/dev/md/openqa"

```

```

Mär 11 15:48:19 openqaworker2 openqa-establish-nvme-setup[5013]: mdadm: Cannot get exclusive access to /de
v/md/openqa:Perhaps a running process, mounted filesystem or active volume group?
Mär 11 15:48:19 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, st
atus=1/FAILURE
Mär 11 15:48:19 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 11 15:48:19 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 11 15:48:19 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 11 15:49:39 openqaworker2 openqa-establish-nvme-setup[5118]: Stopping current RAID "/dev/md/openqa"
Mär 11 15:49:39 openqaworker2 openqa-establish-nvme-setup[5118]: mdadm: Cannot get exclusive access to /de
v/md/openqa:Perhaps a running process, mounted filesystem or active volume group?
Mär 11 15:49:39 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 11 15:49:39 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, st
atus=1/FAILURE
Mär 11 15:49:39 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 11 15:49:39 openqaworker2 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 11 15:49:39 openqaworker2 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.

```

So here we just get mdadm: cannot open /dev/nvme0n1: Device or resource busy. We've seen Device or resource busy before but only after starting the service a 2nd time when the filesystem has already been mounted. Then it is actually expected but now I'm wondering why the device was considered busy on the first attempt. It then actually worked on a further attempt.

The log shows that I've also restarted the service a few more times. As before this doesn't work. With the SR we no longer get Device or resource busy in that case because I don't suppress the error from the attempt to stop the RAID anymore. So the error preventing the stopping (mdadm: Cannot get exclusive access to /dev/md/openqa:Perhaps a running process, mounted filesystem or active volume group?) is logged instead.

To shed some light on this behavior I've added further debugging output (current block devs and mounts) but I couldn't reproduce Device or resource busy anymore.

#24 - 2021-03-15 13:03 - mkittler

After further tests on openqaworker13 the SR has been merged. I can now try to implement a re-try.

#25 - 2021-03-15 18:50 - mkittler

- Status changed from In Progress to Feedback

SR for adding a retry: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/466

I've been rebooting openqaworker2 a few times with this change. Unfortunately it always worked right away so I couldn't really test the retry logic in production. However, at least this change doesn't break what was working before.

#26 - 2021-03-16 10:39 - mkittler

Observed the problem on openqaworker3 as well. Of course there the SR wasn't applied so it didn't help much with testing. However, now we know that this isn't a machine-specific problem. Strangely, openqaworker3 did *not* end up in the emergency shell despite the same service failing:

```

martchus@openqaworker3:/srv/salt> sudo journalctl --boot -u openqa_nvme_format.service
-- Logs begin at Fri 2021-03-05 06:38:36 CET, end at Tue 2021-03-16 11:32:32 CET. --
Mär 14 03:34:38 openqaworker3 sh[734]: └─sda3      8:3    0 920.8G  0 part /
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 14 03:34:38 openqaworker3 sh[734]:          level=raid0 devices=2 ctime=Sun Feb 28 03:34:44 2021
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or
Mär 14 03:34:38 openqaworker3 sh[734]:          meaningless after creating array
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: /dev/nvme1n1 appears to be part of a raid array:
Mär 14 03:34:38 openqaworker3 sh[734]:          level=raid0 devices=2 ctime=Sun Feb 28 03:34:44 2021
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: partition table exists on /dev/nvme1n1 but will be lost or
Mär 14 03:34:38 openqaworker3 sh[734]:          meaningless after creating array
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: Defaulting to version 1.2 metadata
Mär 14 03:34:38 openqaworker3 sh[734]: mdadm: array /dev/md/openqa started.
Mär 14 03:34:43 openqaworker3 sh[734]: mdadm: timeout waiting for /dev/md/openqa
Mär 14 03:34:43 openqaworker3 sh[848]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 14 03:34:43 openqaworker3 sh[849]: ARRAY /dev/md127 metadata=1.2 name=openqaworker3:openqa UUID=83295cf5:4
a378c19:568da609:be95fbae
Mär 14 03:34:43 openqaworker3 mkfs.ext2[852]: mke2fs 1.43.8 (1-Jan-2018)
Mär 14 03:34:43 openqaworker3 mkfs.ext2[852]: The file /dev/md/openqa does not exist and no size was specified
.
Mär 14 03:34:43 openqaworker3 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=1/FAILURE
Mär 14 03:34:43 openqaworker3 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 14 03:34:43 openqaworker3 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 14 03:34:43 openqaworker3 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 14 03:34:47 openqaworker3 systemd[1]: Starting Setup NVMe before mounting it...
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 14 03:34:47 openqaworker3 sh[948]:          level=raid0 devices=2 ctime=Sun Mar 14 03:34:38 2021
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or

```

```

Mär 14 03:34:47 openqaworker3 sh[948]:      meaningless after creating array
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: /dev/nvme1n1 appears to be part of a raid array:
Mär 14 03:34:47 openqaworker3 sh[948]:      level=raid0 devices=2 ctime=Sun Mar 14 03:34:38 2021
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: partition table exists on /dev/nvme1n1 but will be lost or
Mär 14 03:34:47 openqaworker3 sh[948]:      meaningless after creating array
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: Defaulting to version 1.2 metadata
Mär 14 03:34:47 openqaworker3 sh[948]: mdadm: array /dev/md/openqa started.
Mär 14 03:34:47 openqaworker3 sh[1008]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 14 03:34:47 openqaworker3 sh[1010]: ARRAY /dev/md/openqa metadata=1.2 name=openqaworker3:openqa UUID=f8431
0ca:70b3638a:422dd89c:3915a06e
Mär 14 03:34:47 openqaworker3 mkfs.ext2[1025]: mke2fs 1.43.8 (1-Jan-2018)
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: [182B blob data]
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: Creating filesystem with 195289600 4k blocks and 48824320 inode
s
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: Filesystem UUID: 65b7393e-97d5-4362-acc0-204d0fc13c1b
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: Superblock backups stored on blocks:
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]:      32768, 98304, 163840, 229376, 294912, 819200, 884736, 1
605632, 2654208,
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]:      4096000, 7962624, 11239424, 20480000, 23887872, 7166361
6, 78675968,
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]:      102400000
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: [65B blob data]
Mär 14 03:34:54 openqaworker3 mkfs.ext2[1025]: [134B blob data]
Mär 14 03:35:09 openqaworker3 mkfs.ext2[1025]: [99B blob data]
Mär 14 03:35:09 openqaworker3 systemd[1]: Started Setup NVMe before mounting it.
Mär 15 06:51:36 openqaworker3 systemd[1]: Starting Setup NVMe before mounting it...
Mär 15 06:51:36 openqaworker3 sh[1131]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
Mär 15 06:51:36 openqaworker3 sh[1131]: mdadm: cannot open /dev/nvme0n1p3: No such file or directory
Mär 15 06:51:37 openqaworker3 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=2/INVALIDARGUMENT
Mär 15 06:51:37 openqaworker3 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 15 06:51:37 openqaworker3 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 15 06:51:37 openqaworker3 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.

```

Or did somebody recover it on Monday after reading the [Alerting] openqaworker3: host up alert from Sunday?

#27 - 2021-03-22 13:11 - mkittler

- File `openqa_nvme_format.service.boot.openqaworker3.log` added
- File `openqa_nvme_format.service.boot.openqaworker8.log` added
- File `openqa_nvme_format.service.boot.openqaworker9.log` added

The "Device or resource busy" problem happened today on openqaworker{3,8,9}.

On openqaworker{8,9} it helped to simply continue from the emergency shell. On openqaworker3 I needed to reboot for `openqa_nvme_format.service` so succeed. I've attached the output of `journalctl --no-pager --boot -u openqa_nvme_format.service > /tmp/openqa_nvme_format.service.boot.openqaworkerX.log`.

#28 - 2021-03-23 09:28 - dzedro

openqaworker2.suse.de is down
update: I just pressed ctrl-d to boot the machine

#29 - 2021-03-23 16:11 - mkittler

Relevant log for openqaworker2:

```

-- Logs begin at Wed 2021-03-10 13:11:14 CET, end at Tue 2021-03-23 17:07:51 CET. --
Mär 21 03:35:08 openqaworker2 systemd[1]: Starting Setup NVMe before mounting it...
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: Current mount points (printed for debugging pu
rposes):
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: sysfs on /sys type sysfs (rw,nosuid,nodev,noex
ec,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: proc on /proc type proc (rw,nosuid,nodev,noexe
c,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: devtmpfs on /dev type devtmpfs (rw,nosuid,size
=131889892k,nr_inodes=32972473,mode=755)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: securityfs on /sys/kernel/security type securi
tyfs (rw,nosuid,nodev,noexec,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: devpts on /dev/pts type devpts (rw,nosuid,noex
ec,relatime,gid=5,mode=620,ptmxmode=000)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: tmpfs on /run type tmpfs (rw,nosuid,nodev,mode
=755)

```

Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/unified type cgroup2 (rw,nosuid,nodev,noexec,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,xattr,name=systemd)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: none on /sys/fs/bpf type bpf (rw,nosuid,nodev,noexec,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,perf_event)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/rdma type cgroup (rw,nosuid,nodev,noexec,relatime,rdma)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup (rw,nosuid,nodev,noexec,relatime,net_cls,net_prio)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,cpuset)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup (rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,freezer)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,memory)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/blkio type cgroup (rw,nosuid,nodev,noexec,relatime,blkio)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,devices)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: cgroup on /sys/fs/cgroup/pids type cgroup (rw,nosuid,nodev,noexec,relatime,pids)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on / type btrfs (rw,noatime,space_cache,subvolid=259,subvol=/@/.snapshots/1/snapshot)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: debugfs on /sys/kernel/debug type debugfs (rw,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: mqueue on /dev/mqueue type mqueue (rw,relatime)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: hugetlbf on /dev/hugepages type hugetlbf (rw,relatime,pagesize=2M)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: systemd-1 on /proc/sys/fs/binfmt_misc type autofs (rw,relatime,fd=32,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=32382)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/mailman type btrfs (rw,noatime,space_cache,subvolid=269,subvol=/@/var/lib/mailman)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/crash type btrfs (rw,noatime,space_cache,subvolid=267,subvol=/@/var/crash)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/pgsql type btrfs (rw,noatime,space_cache,subvolid=273,subvol=/@/var/lib/pgsql)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /usr/local type btrfs (rw,noatime,space_cache,subvolid=266,subvol=/@/usr/local)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /tmp type btrfs (rw,noatime,space_cache,subvolid=265,subvol=/@/tmp)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/opt type btrfs (rw,noatime,space_cache,subvolid=275,subvol=/@/var/opt)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/mysql type btrfs (rw,noatime,space_cache,subvolid=271,subvol=/@/var/lib/mysql)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /.snapshots type btrfs (rw,noatime,space_cache,subvolid=258,subvol=/@/.snapshots)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /boot/grub2/i386-pc type btrfs (rw,noatime,space_cache,subvolid=260,subvol=/@/boot/grub2/i386-pc)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /opt type btrfs (rw,noatime,space_cache,subvolid=263,subvol=/@/opt)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/libvirt/images type btrfs (rw,noatime,space_cache,subvolid=268,subvol=/@/var/lib/libvirt/images)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /home type btrfs (rw,noatime,space_cache,subvolid=262,subvol=/@/home)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/tmp type btrfs (rw,noatime,space_cache,subvolid=277,subvol=/@/var/tmp)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/machines type btrfs (rw,noatime,space_cache,subvolid=506,subvol=/@/var/lib/machines)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/mariadb type btrfs (rw,noatime,space_cache,subvolid=270,subvol=/@/var/lib/mariadb)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/log type btrfs (rw,noatime,space_cache,subvolid=274,subvol=/@/var/log)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/lib/named type btrfs (rw,noatime,space_cache,subvolid=272,subvol=/@/var/lib/named)

```

ime,space_cache,subvolid=272,subvol=/@/var/lib/named)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /var/spool type btrfs (rw,noatime,
space_cache,subvolid=276,subvol=/@/var/spool)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /srv type btrfs (rw,noatime,space_
cache,subvolid=264,subvol=/@/srv)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: /dev/md1 on /boot/grub2/x86_64-efi type btrfs
(rw,noatime,space_cache,subvolid=261,subvol=/@/boot/grub2/x86_64-efi)
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: Present block devices (printed for debugging p
urposes):
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: NAME          MAJ:MIN RM   SIZE RO TYPE  MOUNTPO
INT
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: sda          8:0      0 931.5G 0 disk
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └sda1       8:1      0    1G 0 part
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: | └md0        9:0      0    2G 0 raid0
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └sda2       8:2      0 930.5G 0 part
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └md1        9:1      0 930.5G 0 raid1 /
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: sdb          8:16     0 931.5G 0 disk
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └sdb1       8:17     0    1G 0 part
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: | └md0        9:0      0    2G 0 raid0
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └sdb2       8:18     0 930.5G 0 part
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └md1        9:1      0 930.5G 0 raid1 /
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: nvme1n1     259:0    0 372.6G 0 disk
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: nvme0n1     259:1    0 372.6G 0 disk
Mär 21 03:35:08 openqaworker2 openqa-establish-nvme-setup[969]: └nvme0n1p1 259:2    0 372.6G 0 part
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: Creating RAID0 "/dev/md/openqa" on: /dev/nvme0
n1 /dev/nvme1n1
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: /dev/nvme0n1 appears to be part of a ra
id array:
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]:          level=raid0 devices=2 ctime=Mon Mar 15
18:55:33 2021
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: partition table exists on /dev/nvme0n1
but will be lost or
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]:          meaningless after creating array
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: /dev/nvme1n1 appears to be part of a ra
id array:
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]:          level=raid0 devices=2 ctime=Mon Mar 15
18:55:33 2021
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: partition table exists on /dev/nvme1n1
but will be lost or
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]:          meaningless after creating array
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: Defaulting to version 1.2 metadata
Mär 21 03:35:10 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: array /dev/md/openqa started.
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: timeout waiting for /dev/md/openqa
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: Trying RAID0 creation again after timeout (att
empt 2 of 4)
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: mdadm: cannot open /dev/nvme0n1: Device or res
ource busy
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: Status for RAID0 "/dev/md/openqa"
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: ARRAY /dev/md127 metadata=1.2 name=openqaworke
r2:openqa UUID=eb901be8:9eb9682e:c5da919c:89125c6d
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: Creating ext2 filesystem on RAID0 "/dev/md/ope
nqa"
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: mke2fs 1.43.8 (1-Jan-2018)
Mär 21 03:35:13 openqaworker2 openqa-establish-nvme-setup[969]: The file /dev/md/openqa does not exist and no
size was specified.
Mär 21 03:35:13 openqaworker2 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=1/FAILURE
Mär 21 03:35:13 openqaworker2 systemd[1]: Failed to start Setup NVMe before mounting it.

```

#30 - 2021-03-29 09:06 - mkittler

This happened again yesterday on openqaworker8 and openqaworker9. Just pressing Ctrl+D helped there as well.

```

-- Logs begin at Tue 2021-03-16 17:17:48 CET, end at Mon 2021-03-29 10:58:57 CEST. --
Mär 28 03:34:19 openqaworker9 systemd[1]: Starting Setup NVMe before mounting it...
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: Current mount points (printed for debugging p
urposes):
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: sysfs on /sys type sysfs (rw,nosuid,nodev,noe
xec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: proc on /proc type proc (rw,nosuid,nodev,noex
ec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: devtmpfs on /dev type devtmpfs (rw,nosuid,siz
e=131888880k,nr_inodes=32972220,mode=755)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: securityfs on /sys/kernel/security type secur

```

ityfs (rw,nosuid,nodev,noexec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: tmpfs on /run type tmpfs (rw,nosuid,nodev,mode=755)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/unified type cgroup2 (rw,nosuid,nodev,noexec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,xattr,name=systemd)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: efivarfs on /sys/firmware/efi/efivars type efivarfs (rw,nosuid,nodev,noexec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: none on /sys/fs/bpf type bpf (rw,nosuid,nodev,noexec,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,devices)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,cpuset)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,perf_event)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/blkio type cgroup (rw,nosuid,nodev,noexec,relatime,blkio)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/rdma type cgroup (rw,nosuid,nodev,noexec,relatime,rdma)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,memory)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup (rw,nosuid,nodev,noexec,relatime,net_cls,net_prio)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup (rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,freezer)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: cgroup on /sys/fs/cgroup/pids type cgroup (rw,nosuid,nodev,noexec,relatime,pids)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on / type btrfs (rw,relatime,ssd,space_cache,subvolid=259,subvol=/@/.snapshots/1/snapshot)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: mqueue on /dev/mqueue type mqueue (rw,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,pagesize=2M)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: debugfs on /sys/kernel/debug type debugfs (rw,relatime)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: systemd-1 on /proc/sys/fs/binfmt_misc type autofs (rw,relatime,fd=36,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=23612)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /boot/grub2/i386-pc type btrfs (rw,relatime,ssd,space_cache,subvolid=260,subvol=/@/boot/grub2/i386-pc)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/log type btrfs (rw,relatime,ssd,space_cache,subvolid=275,subvol=/@/var/log)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /srv type btrfs (rw,relatime,ssd,space_cache,subvolid=263,subvol=/@/srv)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /.snapshots type btrfs (rw,relatime,ssd,space_cache,subvolid=258,subvol=/@/.snapshots)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/mailman type btrfs (rw,relatime,ssd,space_cache,subvolid=270,subvol=/@/var/lib/mailman)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p1 on /boot/efi type vfat (rw,relatime,mask=0002,dmask=0002,allow_utime=0020,codepage=437,iocharset=iso8859-1,shortname=mixed,utf8,errors=remount-ro)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/mysql type btrfs (rw,relatime,ssd,space_cache,subvolid=272,subvol=/@/var/lib/mysql)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/tmp type btrfs (rw,relatime,ssd,space_cache,subvolid=278,subvol=/@/var/tmp)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/machines type btrfs (rw,relatime,ssd,space_cache,subvolid=269,subvol=/@/var/lib/machines)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /tmp type btrfs (rw,relatime,ssd,space_cache,subvolid=264,subvol=/@/tmp)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/libvirt/images type btrfs (rw,relatime,ssd,space_cache,subvolid=268,subvol=/@/var/lib/libvirt/images)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/crash type btrfs (rw,r

```

elotime,ssd,space_cache,subvolid=267,subvol=/@/var/crash)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /opt type btrfs (rw,relatime,ssd,space_cache,subvolid=262,subvol=/@/opt)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/named type btrfs (rw,relatime,ssd,space_cache,subvolid=273,subvol=/@/var/lib/named)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/cache type btrfs (rw,relatime,ssd,space_cache,subvolid=266,subvol=/@/var/cache)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/pgsqli type btrfs (rw,relatime,ssd,space_cache,subvolid=274,subvol=/@/var/lib/pgsqli)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /usr/local type btrfs (rw,relatime,ssd,space_cache,subvolid=265,subvol=/@/usr/local)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/lib/mariadb type btrfs (rw,relatime,ssd,space_cache,subvolid=271,subvol=/@/var/lib/mariadb)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /boot/grub2/x86_64-efi type btrfs (rw,relatime,ssd,space_cache,subvolid=261,subvol=/@/boot/grub2/x86_64-efi)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/spool type btrfs (rw,relatime,ssd,space_cache,subvolid=277,subvol=/@/var/spool)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: /dev/nvme0n1p2 on /var/opt type btrfs (rw,relatime,ssd,space_cache,subvolid=276,subvol=/@/var/opt)
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: Present block devices (printed for debugging purposes):
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: NAME          MAJ:MIN RM   SIZE RO TYPE  MOUNTPOINT
OINT
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: nvme0n1      259:0    0 745.2G  0 disk
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: |-nvme0n1p1 259:1    0  156M  0 part  /boot/efi
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: |-nvme0n1p2 259:2    0  100G  0 part  /
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: |-nvme0n1p3 259:3    0 643.1G  0 part
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: |  md127      9:127   0 642.9G  0 raid0
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: Creating RAID0 "/dev/md/openqa" on: /dev/nvme0n1
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: Stopping current RAID "/dev/md/openqa"
Mär 28 03:34:19 openqaworker9 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status=1/FAILURE
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: mdadm: stopped /dev/md/openqa
Mär 28 03:34:19 openqaworker9 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
Mär 28 03:34:19 openqaworker9 openqa-establish-nvme-setup[1283]: Status for RAID0 "/dev/md/openqa"
Mär 28 03:34:19 openqaworker9 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 28 03:34:19 openqaworker9 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.
Mär 29 10:52:07 openqaworker9 systemd[1]: Starting Setup NVMe before mounting it...
... then, after Ctrl-D in the rescue shell the same mount/lsblk output but everything works fine.

```

So the Device or resource busy problem can happen in different places and a retry helps. It looks *not* like it is caused by the partition already being mounted in these cases.

#31 - 2021-03-29 13:42 - mkittler

SR: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/468

Now, I could also reproduce a case where the re-try was triggered and actually helped (in all my tests so far it just succeeded on the first try):

```

Mär 29 15:08:05 openqaworker3 openqa-establish-nvme-setup[907]: Creating RAID0 "/dev/md/openqa" on: /dev/nvme0n1 /dev/nvme1n1
Mär 29 15:08:05 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: cannot open /dev/nvme0n1: Device or resource busy
Mär 29 15:08:05 openqaworker3 openqa-establish-nvme-setup[907]: Waiting 10 seconds before trying again after failing due to busy device.
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: Trying RAID0 creation again after timeout (attempt 2 of 10)
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: Stopping current RAID "/dev/md/openqa"
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: stopped /dev/md/openqa
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: /dev/nvme0n1 appears to be part of a raid array:
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: level=raid0 devices=2 ctime=Mon Mar 29 15:01:13 2021
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: partition table exists on /dev/nvme0n1 but will be lost or
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: meaningless after creating array
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: /dev/nvme1n1 appears to be part of a raid array:
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: level=raid0 devices=2 ctime=Mon Mar 29 15:01:13 2021
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: partition table exists on /dev/nvme1n1

```



```

but will be lost or
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]:      meaningless after creating array
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: Defaulting to version 1.2 metadata
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mdadm: array /dev/md/openqa started.
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: Status for RAID0 "/dev/md/openqa"
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: md127 : active raid0 nvme1n1[1] nvme0n1[0]
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: ARRAY /dev/md/openqa metadata=1.2 name=openqaw
orker3:openqa UUID=85b8953d:f6bab579:cec0eb6f:ea8855ce
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: Creating ext2 filesystem on RAID0 "/dev/md/ope
nqa"
Mär 29 15:08:11 openqaworker3 openqa-establish-nvme-setup[907]: mke2fs 1.43.8 (1-Jan-2018)
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]: [182B blob data]
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]: Creating filesystem with 195289600 4k blocks a
nd 48824320 inodes
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]: Filesystem UUID: f3c8ac1a-f4ea-4e68-93e7-77a6c
b52f05c
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]: Superblock backups stored on blocks:
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]:      32768, 98304, 163840, 229376, 294912,
819200, 884736, 1605632, 2654208,
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]:      4096000, 7962624, 11239424, 20480000,
23887872, 71663616, 78675968,
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]:      102400000
Mär 29 15:08:13 openqaworker3 openqa-establish-nvme-setup[907]: [65B blob data]
Mär 29 15:08:17 openqaworker3 openqa-establish-nvme-setup[907]: [134B blob data]
Mär 29 15:08:32 openqaworker3 openqa-establish-nvme-setup[907]: [99B blob data]
Mär 29 15:08:32 openqaworker3 systemd[1]: Started Setup NVMe before mounting it.

```

#32 - 2021-03-29 15:52 - mkittler

I've been rebooting openqaworker3 lots of times and most of them it just succeeded. However, now I encountered yet another type of error:

```

Mär 29 17:07:08 openqaworker3 openqa-establish-nvme-setup[899]: Creating RAID0 "/dev/md/openqa" on: /dev/nvme0
n1 /dev/nvme1n1
Mär 29 17:07:08 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: cannot open /dev/nvme0n1: Device or res
ource busy
Mär 29 17:07:08 openqaworker3 openqa-establish-nvme-setup[899]: Waiting 10 seconds before trying again after f
ailing due to busy device.
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: Trying RAID0 creation again after timeout (att
empt 2 of 10)
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: Stopping current RAID "/dev/md/openqa"
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: stopped /dev/md/openqa
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: /dev/nvme0n1 appears to be part of a ra
id array:
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]:      level=raid0 devices=2 ctime=Mon Mar 29
16:55:14 2021
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: partition table exists on /dev/nvme0n1
but will be lost or
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]:      meaningless after creating array
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: /dev/nvme1n1 appears to be part of a ra
id array:
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]:      level=raid0 devices=2 ctime=Mon Mar 29
16:55:14 2021
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: partition table exists on /dev/nvme1n1
but will be lost or
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]:      meaningless after creating array
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: mdadm: unexpected failure opening /dev/md127
Mär 29 17:07:17 openqaworker3 openqa-establish-nvme-setup[899]: Unable to create RAID, mdadm returned with non
-zero code
Mär 29 17:07:17 openqaworker3 systemd[1]: openqa_nvme_format.service: Main process exited, code=exited, status
=1/FAILURE
Mär 29 17:07:17 openqaworker3 systemd[1]: Failed to start Setup NVMe before mounting it.
Mär 29 17:07:17 openqaworker3 systemd[1]: openqa_nvme_format.service: Unit entered failed state.
Mär 29 17:07:17 openqaworker3 systemd[1]: openqa_nvme_format.service: Failed with result 'exit-code'.

```

Since the retry from https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/468 is not applied blindly this was treated as fatal error and the worker ended up in the rescue shell. Considering that simply retrying (without rebooting) helped in this error case as well it would be tempting to enable the retry here as well. However, the error message "unexpected failure" is quite unspecific so we would possibly retry in cases where it doesn't make sense.

#33 - 2021-04-06 09:38 - mkittler

SR for problems seen on openqaworker 8 and 9: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/474

#34 - 2021-04-08 15:02 - okurz

- Status changed from Feedback to In Progress

- Assignee changed from mkittler to okurz

debugging session okurz+mkittler. A previous commit introduced a problem. With [#78010](#) resolved I suggest we continue here and crosscheck stability of openqaworker2. I am taking openqaworker2 offline and will conduct a thorough check there:

```
systemctl mask --now telegraf salt-minion openqa-worker-auto-restart@{1..28}
```

and from another machine (lord.arch):

```
for run in {01..30}; do for host in openqaworker2; 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
```

#35 - 2021-04-08 19:43 - okurz

- Status changed from In Progress to Resolved

- Assignee changed from okurz to mkittler

all reboots successful. both ACs completed.

Files

Screenshot_2021-01-24_22-09-45.png	42.6 KB	2021-01-25	dzedro
Screenshot_2021-01-25_14-43-16.png	91.5 KB	2021-01-25	dzedro
openqa_nvme_format.service.boot.openqaworker3.log	37.8 KB	2021-03-22	mkittler
openqa_nvme_format.service.boot.openqaworker8.log	21.1 KB	2021-03-22	mkittler
openqa_nvme_format.service.boot.openqaworker9.log	20.8 KB	2021-03-22	mkittler