

## openQA Infrastructure - action #69523

### lessons learned: osd did not come up after reboot 2020-08-02

2020-08-03 06:52 - okurz

<b>Status:</b>	Resolved	<b>Start date:</b>	2020-08-03
<b>Priority:</b>	Urgent	<b>Due date:</b>	
<b>Assignee:</b>	okurz	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Ready		

#### Description

### Observation

After [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/314](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/314) enabled automatic reboot of osd itself the VM did not come up from reboot on 2020-08-02 . okurz reported <https://infra.nue.suse.com/SelfService/Display.html?id=175461> which bmwiedemann could resolve early 2020-08-03 . okurz logged in as root with the SSH key from "backup-vm" as normal user login was not working. Many services were not running. The mounted partition were not in order:

```
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       9.6G  5.8G  3.4G  63% /
/dev/vdc        5.0T  4.5T  605G  89% /srv
/dev/vdd        5.0T  3.4T  1.7T  68% /assets
```

missing /home and /results and also /srv should not be 5TB.

### Problem

During the course of the last two years likely the former volume "vdb" was removed during runtime of the VM at the time when coolo and EngInfra moved assets and results to two new, separated volumes. However as /etc/fstab relies on the order of detected devices after former vdb vanished on the next bootup the existing partitions are now identified differently and were assigned to incorrect mount points and also former "vde" was missing.

### Solution

okurz fixed this with defining mount points in /etc/fstab using UUIDs:

```
# cat /etc/fstab
devpts /dev/pts          devpts mode=0620,gid=5 0 0
proc   /proc                proc   defaults                0 0
sysfs  /sys                  sysfs  noauto                  0 0
debugfs /sys/kernel/debug    debugfs noauto                  0 0
usbfs  /proc/bus/usb         usbfs  noauto                  0 0
tmpfs  /run                  tmpfs  noauto                  0 0
# 7116dc72-ebc8-4b21-8847-b9f31dc95229 -> vda1
/dev/vda1 / ext3 defaults 1 1
# 2e55520d-2b90-4100-8892-025c5f4c9949 -> vda2
/dev/vda2 swap swap defaults 0 0
# 6c8044d6-5497-4db6-9714-89b76268121e -> vdb
UUID=6c8044d6-5497-4db6-9714-89b76268121e /srv xfs defaults,logbsize=256k,noatime,nodiratime 1 2
/srv/PSQL10 /var/lib/pgsql none bind 0 0
# 3f003a69-c51e-4d79-8b83-906e7918bac4 -> vdc
UUID=3f003a69-c51e-4d79-8b83-906e7918bac4 /assets xfs defaults,logbsize=256k,noatime,nodiratime 1 2
/assets /var/lib/openqa/share none bind 0 0
# 51d504aa-6f46-4b89-bcd9-b6cea7b8b755 -> vdd
UUID=51d504aa-6f46-4b89-bcd9-b6cea7b8b755 /results xfs defaults,logbsize=256k,noatime,nodiratime 1 2
/results /var/lib/openqa none bind 0 0
/srv/homes.img /home ext4 defaults 1 1
```

As an alternative labels could be used. However only "assets" is currently available with a label. Filesystem labels can be set with

tune2fs -L \$LABEL /dev/vd\$i for ext2/3/4 or for xfs xfs\_admin -L \$LABEL /dev/vd\$i which however needs unmounted volumes so something left for later.

## History

### #1 - 2020-08-03 07:38 - okurz

- Description updated

- Priority changed from Immediate to Urgent

Situation was resolved. All services on osd operational again. See description for main problem. Additionally:

- systemctl default
- systemctl reset-failed
- rm /run/nologin to prevent "System is booting up. See pam\_nologin(8)" for non-root ssh login attempts
- Triggered manually systemctl start openqa-enqueue-asset-cleanup.timer openqa-enqueue-audit-event-cleanup.timer openqa-enqueue-bug-cleanup.timer openqa-enqueue-result-cleanup.timer
- cleaned up some data on partitions, e.g. /assets/log/ /assets/PSQL10/
- sudo systemctl restart nfs-server
- sudo salt -l error -C 'G@roles:worker' cmd.run 'mount -a' to repair the mount of var-lib-openqa-share.mount
- New jobs were not scheduled because of

```
okurz@openqa:~> sudo systemctl status openqa-scheduler
```

```
● openqa-scheduler.service - The openQA Scheduler
   Loaded: loaded (/usr/lib/systemd/system/openqa-scheduler.service; enabled; vendor preset: disabled)
   Drop-In: /etc/systemd/system/openqa-scheduler.service.d
            └─override.conf
   Active: active (running) since Mon 2020-08-03 08:46:31 CEST; 45min ago
   Main PID: 31292 (openqa-schedule)
     Tasks: 1
    CGroup: /system.slice/openqa-scheduler.service
            └─31292 /usr/bin/perl /usr/share/openqa/script/openqa-scheduler daemon -m production
```

```
Aug 03 08:46:31 openqa systemd[1]: Started The openQA Scheduler.
```

```
Aug 03 08:46:34 openqa openqa-scheduler-daemon[31292]: [2020-08-03 08:46:34.05381] [31292] [warn] Deprecated use of config key '[audit]: blacklist'. Use '[audit]: blocklist' instead
```

```
Aug 03 08:46:34 openqa openqa-scheduler-daemon[31292]: Mojo::Reactor::Poll: Timer failed: Can't open database lock file /var/lib/openqa/db/db.lock! at /usr/share/openqa/script/../lib/OpenQA/Schema.pm line 87.
```

Fixed by restart of the service but IMHO the service should not stay running with Mojo::Reactor::Poll: Timer failed: Can't open database lock file /var/lib/openqa/db/db.lock! at /usr/share/openqa/script/../lib/OpenQA/Schema.pm line 87. Recorded now in [#65271#note-24](#)

- sudo salt -l error -C 'G@roles:worker' cmd.run 'systemctl is-system-running' is good now, 144 jobs currently running.
- Triggered explicitly <https://gitlab.suse.de/openqa/auto-review/-/pipelines/73065>
- Monitoring <https://gitlab.suse.de/openqa/openqa-review/-/jobs/239542>

### #2 - 2020-08-03 13:08 - okurz

- Status changed from In Progress to Resolved

These jobs look fine so far as well (except for already existing problems fixed elsewhere).