

## openQA Infrastructure - action #93683

### osd-deployment failed due to storage.qa.suse.de not reachable by salt

2021-06-09 06:38 - okurz

|                              |                                  |
|------------------------------|----------------------------------|
| <b>Status:</b> Resolved      | <b>Start date:</b> 2021-06-09    |
| <b>Priority:</b> Urgent      | <b>Due date:</b>                 |
| <b>Assignee:</b> nicksinger  | <b>% Done:</b> 0%                |
| <b>Category:</b>             | <b>Estimated time:</b> 0.00 hour |
| <b>Target version:</b> Ready |                                  |

**Description**

### Observation

<https://gitlab.suse.de/openqa/osd-deployment/-/jobs/453045> shows

```
storage.qa.suse.de:
  Minion did not return. [Not connected]
```

ipmi-ipmi.storage.qa sol activate shows

```
storage login: root
Password:
Have a lot of fun...
-bash-4.4#
```

so no proper PS1. And there is no systemd running.

### Acceptance criteria

- **AC1:** storage.qa is back
- **AC1:** osd deployment continues after storage.qa is back

### Suggestions

- Remove storage.qa from salt control with ssh osd 'sudo salt-key -y -d storage.qa.suse.de'
- Try to reboot storage.qa and see what happens
- Check reboot stability of storage.qa

### Out of scope

- Monitoring for storage.qa: [#91779](#)

### Rollback

- Add storage.qa back to salt ssh osd 'sudo salt-key -y -a storage.qa.suse.de'

**Related issues:**

|  |                 |                   |
|--|-----------------|-------------------|
| Related to openQA Infrastructure - action #90629: administration of the new "... | <b>Resolved</b> | <b>2020-08-04</b> |
| Related to openQA Infrastructure - action #69577: Handle installation of the ... | <b>Resolved</b> | <b>2020-08-04</b> |
| Related to openQA Infrastructure - action #66709: Storage server for OSD and ... | <b>Resolved</b> | <b>2020-05-12</b> |

### History

#### #1 - 2021-06-09 06:39 - okurz

I triggered a reboot with echo b >/proc/sysrq-trigger and now the machine is stuck in a grub command line

#### #2 - 2021-06-09 06:44 - okurz

- Description updated

As storage.qa.suse.de is now stuck in grub prompt I did ssh osd 'sudo salt-key -y -d storage.qa.suse.de' to be able to continue with the deployment:

**#3 - 2021-06-09 16:01 - cdywan**

Since I made the mistake of thinking Oli was investigating this issue earlier: The above is just the MR to drop storage.qa.suse.de from salt.

- A naive ssh storage.qa.suse.de appears to time out here, so *something* is responding but not really
- IPMI command from [workerconf.sls](#) connects successfully
- exiting *grub* landed me in UEFI
- After waiting a little while for "Checking Media Presence" to do something I am back in *grub*
- *local* puts me back in the *grub* command line
- *normal* doesn't work

Do we have any documentation on what should be booted here? Devices used?

**#4 - 2021-06-09 19:11 - okurz**

cdywan wrote:

Since I made the mistake of thinking Oli was investigating this issue earlier: The above is just the MR to drop storage.qa.suse.de from salt.

It was not a "MR", just a command

Do we have any documentation on what should be booted here? Devices used?

Search for previous tickets about the host, e.g. look for "storage". nicksinger conducted the installation of the host. I assume no profile for automated installation exists, e.g. no autoyast profile, just a manual installation. Although I am thinking for all newly installed machines we should aim for a completely automated installation from autoyast.

**#5 - 2021-06-10 07:45 - nicksinger**

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

- Assignee set to *nicksinger*

The machine is in a really strange state. I've booted a rescue media and chrooted into the rootfs (on nvme0n1p2) where many things seem to be missing. E.g. zypper is missing (the whole /usr/bin folder is missing) and also snapper is not present. However the files are present in snapshots. I will check if I can maybe restore one of the snapshots and investigate why blew up the OS there

**#6 - 2021-06-10 11:58 - nicksinger**

After playing around with the btrfs on there I realized that the system is mounting snapshot 1 as default (described as "first root filesystem" by snapper). This snapshot really contains almost nothing. I was able to manually mount snapshot 33 (the most recent one) with `btrfs mount /dev/nvme0n1p2 /mnt/mychroot -o subvol=@/.snapshots/33/snapshot` and chroot into that snapshot. It was ro (as snapshots always are) but snapper was installed in there so I could do a snapper rollback 33 which created snapshot 35 (writable copy of #33) and set it as default. After a `mount -a` inside the chroot I executed `grub2-mkconfig -o /boot/grub2/grub.cfg` just to make sure the correct grub files are generated again. With these changes I was able to reboot the machine again normally. However, for some reason `polkit.service` fails to come up now which (I assume) results in no network. Investigating further if this can easily be fixed now

**#7 - 2021-06-10 12:16 - nicksinger**

got `polkit` (and `cron`) back and running by manually creating `/var/lib/polkit/` and `/var/spool/cron/`. Did another system upgrade and rebooted 3 times without issues. I'm declaring the machine as stable again but I have absolutely no clue how this could happen without somebody manually deleting files.

**#8 - 2021-06-10 12:22 - okurz**

Hi Nick, this is great. Thank you for the quick reaction and detailed update.

cdywan wrote:

Do we have any documentation on what should be booted here? Devices used?

can you remind us about that context? Should we consider cheap redeploy based on autoyast profiles?

**#9 - 2021-06-10 12:23 - nicksinger**

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

I've added the machine back to salt on OSD and ran an highstate to see if maybe this caused the destruction. But everything went smooth and another reboot showed that the machine can be considered stable for now:

```
openqa:~ # salt-key -y -a storage.qa.suse.de
The following keys are going to be accepted:
Unaccepted Keys:
storage.qa.suse.de
Key for minion storage.qa.suse.de accepted.
openqa:~ # salt 'storage.qa.suse.de' test.ping
storage.qa.suse.de:
    True
openqa:~ # salt 'storage.qa.suse.de' state.highstate
storage.qa.suse.de:

Summary for storage.qa.suse.de
-----
Succeeded: 192
Failed:    0
-----
Total states run:    192
Total run time:     7.146 s
```

#### #10 - 2021-06-10 12:26 - nicksinger

okurz wrote:

Hi Nick, this is great. Thank you for the quick reaction and detailed update.

cdywan wrote:

Do we have any documentation on what should be booted here? Devices used?

can you remind us about that context? Should we consider cheap redeploy based on autoyast profiles?

I think it could help even though it masks problems which I don't like. However, instead of autoyast I'd take a look into the yomi project to have an installation based on salt.

#### #11 - 2021-06-10 19:55 - okurz

- Related to action #90629: administration of the new "Storage Server" added

#### #12 - 2021-06-10 19:55 - okurz

- Related to action #69577: Handle installation of the new "Storage Server" added

#### #13 - 2021-06-10 19:56 - okurz

- Related to action #66709: Storage server for OSD and monitoring added

#### #14 - 2021-06-10 19:56 - okurz

nicksinger wrote:

okurz wrote:

Hi Nick, this is great. Thank you for the quick reaction and detailed update.

cdywan wrote:

Do we have any documentation on what should be booted here? Devices used?

can you remind us about that context?

Maybe you missed this question. I linked three tickets for reference.

Should we consider cheap redeploy based on autoyast profiles?

I think it could help even though it masks problems which I don't like.

It should be as simple as <https://w3.nue.suse.com/~okurz/ay-openqa-worker.xml> as the rest is done by salt.

However, instead of autoyast I'd take a look into the yomi project to have an installation based on salt.

yes, I would like that as well.