

## openQA Infrastructure - action #65154

### root partition on osd exceeds alert threshold, 90%, after osd deployment -> apply automatic reboots to OSD machines

2020-04-01 14:43 - okurz

<b>Status:</b>	Resolved	<b>Start date:</b>	2020-04-01
<b>Priority:</b>	High	<b>Due date:</b>	2020-08-04
<b>Assignee:</b>	okurz	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Ready		
<b>Description</b>			
<b>Observation</b>			
<a href="https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&amp;panelId=74&amp;fullscreen&amp;edit&amp;tab=alert&amp;refresh=30s&amp;from=now-7d&amp;to=now">https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&amp;panelId=74&amp;fullscreen&amp;edit&amp;tab=alert&amp;refresh=30s&amp;from=now-7d&amp;to=now</a>			
shows that the osd deployment of today, 2020-04-01, hit the alert threshold for disk usage on / with 90% full.			
<b>Related issues:</b>			
Related to openQA Project - action #69355: [spike] redundant/load-balancing w...		<b>Resolved</b>	<b>2020-07-25</b>

## History

### #1 - 2020-04-01 14:48 - okurz

- Priority changed from Urgent to High

```
du -x -dl -BM
2M    ./bin
1M    ./lost+found
1M    ./selinux
12M   ./lib64
5M    ./sbin
1M    ./mnt
3817M ./usr
44M   ./root
25M   ./etc
212M  ./opt
311M  ./boot
829M  ./var
220M  ./tmp
2883M ./lib
8355M .
```

I guess one problem is too many kernels:

```
# ls /boot/
.vmlinuz-4.12.14-lp151.28.13-default.hmac System.map-4.12.14-lp151.28.20-default config-4.12.14-lp151.28.25-
default initrd-4.12.14-lp151.28.25-default symvers-4.12.14-lp151.28.25-default.gz sysctl.conf-4.12.14-
lp151.28.40-default vmlinux-4.4.159-73-default.gz
.vmlinuz-4.12.14-lp151.28.16-default.hmac System.map-4.12.14-lp151.28.25-default config-4.12.14-lp151.28.32-
default initrd-4.12.14-lp151.28.32-default symvers-4.12.14-lp151.28.32-default.gz sysctl.conf-4.12.14-
lp151.28.44-default vmlinux
.vmlinuz-4.12.14-lp151.28.20-default.hmac System.map-4.12.14-lp151.28.32-default config-4.12.14-lp151.28.36-
default initrd-4.12.14-lp151.28.36-default symvers-4.12.14-lp151.28.36-default.gz sysctl.conf-4.4.159-
73-default vmlinux-4.12.14-lp151.28.13-default
.vmlinuz-4.12.14-lp151.28.25-default.hmac System.map-4.12.14-lp151.28.36-default config-4.12.14-lp151.28.40-
default initrd-4.12.14-lp151.28.40-default symvers-4.12.14-lp151.28.40-default.gz unicode.pf2
.vmlinuz-4.12.14-lp151.28.32-default.hmac System.map-4.12.14-lp151.28.44-default config-4.12.14-lp151.28.44-
default initrd-4.12.14-lp151.28.44-default symvers-4.12.14-lp151.28.44-default.gz vmlinux-4.12.14-lp15
1.28.13-default.gz vmlinux-4.12.14-lp151.28.20-default
.vmlinuz-4.12.14-lp151.28.36-default.hmac System.map-4.12.14-lp151.28.44-default config-4.4.159-73-default
initrd-4.4.159-73-default symvers-4.4.159-73-default.gz vmlinux-4.12.14-lp15
1.28.16-default.gz vmlinux-4.12.14-lp151.28.25-default
.vmlinuz-4.12.14-lp151.28.40-default.hmac System.map-4.4.159-73-default do_purge_kernels
mbrid sysctl.conf-4.12.14-lp151.28.13-default vmlinux-4.12.14-lp15
1.28.20-default.gz vmlinux-4.12.14-lp151.28.32-default
```

```

.vmlinuz-4.12.14-lp151.28.44-default.hmac boot grub2
memtest.bin sysctl.conf-4.12.14-lp151.28.16-default vmlinux-4.12.14-lp15
1.28.25-default.gz vmlinuz-4.12.14-lp151.28.36-default
.vmlinuz-4.4.159-73-default.hmac boot.readme initrd
message sysctl.conf-4.12.14-lp151.28.20-default vmlinux-4.12.14-lp15
1.28.32-default.gz vmlinuz-4.12.14-lp151.28.40-default
0x91f63571 config-4.12.14-lp151.28.13-default initrd-4.12.14-lp151.28.13-
default symvers-4.12.14-lp151.28.13-default.gz sysctl.conf-4.12.14-lp151.28.25-default vmlinux-4.12.14-lp15
1.28.36-default.gz vmlinuz-4.12.14-lp151.28.44-default
System.map-4.12.14-lp151.28.13-default config-4.12.14-lp151.28.16-default initrd-4.12.14-lp151.28.16-
default symvers-4.12.14-lp151.28.16-default.gz sysctl.conf-4.12.14-lp151.28.32-default vmlinux-4.12.14-lp15
1.28.40-default.gz vmlinuz-4.4.159-73-default
System.map-4.12.14-lp151.28.16-default config-4.12.14-lp151.28.20-default initrd-4.12.14-lp151.28.20-
default symvers-4.12.14-lp151.28.20-default.gz sysctl.conf-4.12.14-lp151.28.36-default vmlinux-4.12.14-lp15
1.28.44-default.gz
openqa:/ # uname -a
Linux openqa 4.12.14-lp151.28.13-default #1 SMP Wed Aug 7 07:20:16 UTC 2019 (0c09ad2) x86_64 x86_64 x86_64 GNU
/Linux
openqa:/ # uptime
16:45:24 up 208 days 3:34, 3 users, load average: 8.58, 9.47, 8.80

```

so first I did:

```
zypper rm -u kernel-default-4.4.159-73.1
```

That fixed the alert.

## #2 - 2020-04-01 14:59 - okurz

- Due date set to 2020-05-01

- Status changed from In Progress to Blocked

To my understanding the other kernel versions would only be deleted after a reboot into a newer version. But because we did not reboot for longer the intermediate, "more recent but still older than newest" kernel versions are piling up. We can fix that by rebooting and then upgrading or just for now zypper purge-kernels which I did now. To prevent this in the future I see three options: 1. Run zypper purge-kernels as part of the automatic upgrade, 2. reboot automatically more often ourselves, 3. ask SUSE-IT to handle this as only they have access to the management interface of the VM running openqa.suse.de. The problem with 2. is that in case a reboot fails then we can not recover ourselves.

The cleanup has provided again 2.6GB of free space, currently using 72% on / so original alert problem resolved.

I created <https://infra.nue.suse.com/SelfService/Display.html?id=166761> with the two alternatives of "give us VM management access" or "SUSE IT applies kernel upgrades and reboot periodically".

## #3 - 2020-04-01 15:44 - mkittler

Option 1. seems to be the most straight forward solution at the first look. However, if the currently *running* kernel is not installed on the disk anymore that can lead to problems because additional kernel modules can not be loaded anymore (e.g. you plugging in a mouse or using VPN after the update for the first time since starting the machine might not work). So it would be the best to ensure that the currently running kernel is not purged or we ensure all required kernel modules are loaded in the first place.

## #4 - 2020-04-01 15:50 - coolo

purge-kernels does not remove the running kernel.

## #5 - 2020-04-16 06:38 - okurz

esujskaja suggests: "I still believe, this is a question how QA organizes their working environment. You could stop automatic updates, if needed. You can, with our help, create a full salt record of the desired config and restore machine relatively easy. But defining a policy, how exactly administering QA hosts is on QA team.

We have in plan to provide more autonomy to the users on the VM level - including possibility to create/reboot/restart the instances. But that would be a big project, which we would barely start till autumn, taking all circumstances into account - that is de facto transferring traditional VN infrastructure into a cloud one. So I'd propose still to organize the machine in a stable way or consider possibility to create an immutable record to restore it, if needed. We can create an additional host on morla for you, so you could organize a failover for your tasks."

My answer in the ticket: "Hi, thanks for the nice offers of help and also the possibility of an additional VM. That can be something to investigate in the future. Failover is a good idea. Regarding the upgrades: We already use salt to have a consistent setup of all the machines. My thought is mainly about availability. I would not stop automatic updates but still we need a way to recover a machine when it's down. We had been looking into the SUSE internal cloud infrastructure already years ago but unfortunately the performance was not sufficient (at this time). For the time being if you can't provide access to us and do not want to handle kernel upgrades and reboots within your team then we will still schedule regular updates along with automatic reboots, e.g. every Sunday night. But you can also suggest another timeslot because just in the case the machine does not come up we need help from EngInfra to recover."

I suggest we have scheduled, regular reboots of all machines to activate new kernels and also flush out some other potential problems as well as

ensure we have consistent, reproducible setups. I created [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/301](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/301) to enable rebootmgr but the tests fail for a reason I do not yet understand.

#### #6 - 2020-04-24 11:52 - okurz

- Status changed from Blocked to Workable

#### #7 - 2020-05-16 09:05 - okurz

- Due date changed from 2020-05-01 to 2020-05-30

- Status changed from Workable to Feedback

[https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/301](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/301) enables automatic reboot. For now I have changed that to only apply to workers. This does not fix any problem on osd but we can use that for experimentation. TODO check status after due date and potentially apply the same to osd – unless we have moved osd to a caasp cluster by then with redundant automatically rebooting machines as backend

#### #8 - 2020-05-26 14:42 - okurz

- Subject changed from root partition on osd exceeds alert threshold, 90%, after osd deployment to root partition on osd exceeds alert threshold, 90%, after osd deployment -> apply automatic reboots to OSD machines

- Status changed from Feedback to Workable

I can see that rebootmgr is active on OSD workers but seems to have no effect so far. No worker did a controlled reboot so far. So it seems like I need to trigger the reboots myself, maybe in the automatic update job. Or I try with actually transactional-update even though the systems have a r/w root fs. Experimenting on openqaworker12:

```
zypper -n in transactional-update rebootmgr
systemctl enable transactional-update.timer rebootmgr
systemctl start transactional-update
```

which installed updates (into a btrfs snapshot) and informed rebootmgr. Next reboot is planned by rebootmgr.

Installed updates include e.g. "mosh-1.3.2+20190710-lp151.9.4.x86\_64.rpm", currently installed "...9.3". After reboot the new version is active. Not sure how that works exactly but ok :)

So I guess we should use transactional-update as well? This most likely does work for us because all workers have btrfs root, probably with snapshots enabled, according to salt -l error --no-color -C 'G@roles:worker' cmd.run "findmnt /". For osd itself we need something different though.

#### #9 - 2020-06-12 20:58 - okurz

- Due date deleted (2020-05-30)

- Status changed from Workable to Feedback

Found out we could also use needs-restarting --reboot hint so I suggest we combine the both, "needs-restarting" and "rebootmgrctl". ->

[https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/313](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/313)

EDIT: 2020-06-14: So workers have been triggered for reboot. openqaworker3 and powerqaworker-qam-1 did not come up, grenache-1 has one failed service "systemd-modules-load.service".

-> [#68050](#) and 68053 for the broken machines

I reset the failed service for now. Should check details.

#### #10 - 2020-07-17 16:42 - cdywan

[okurz](#) With the MR merged, do you plan on further steps here?

#### #11 - 2020-07-18 10:07 - okurz

sorry, the last status wasn't clear. The above PR applies automatic reboot to all machines except osd itself whereas the original problem actually only affected osd. [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/314](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/314) is the corresponding new MR for OSD as well which isn't merged yet. The problem is that as we ourselves do not have management access to the OSD machine in case of reboot problems there can be longer unavailability. We had a discussion with Engineering Infrastructure and they do not / can not provide us this access. They could provide us a second VM for redundancy. But before doing a full-blown HA setup we might be better off with a k8s cluster or so.

#### #12 - 2020-07-25 15:17 - okurz

- Related to action #69355: [spike] redundant/load-balancing webui deployments of openQA added

#### #13 - 2020-07-29 07:09 - okurz

- Target version set to Ready

**#14 - 2020-07-29 07:47 - okurz**

- *Due date set to 2020-08-04*

IMHO European summer vacation period is a good time to try out the automatic reboot after updates for the osd VM as well. I merged [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/314](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/314) . Any next update cycle should trigger a reboot next Sunday morning for osd as well. So 2020-08-04 I should check what happened. Similar for openqa-monitor by the way.

**#15 - 2020-08-03 13:05 - okurz**

- *Status changed from Feedback to Resolved*

The last reboot was triggered automatically successfully but did not finish without problems, see [#69523](#) . But we booted the new kernel and /boot shows only "current" and "current-1". So I assume the idea of the original ticket's problem is covered.