

openQA Infrastructure - action #76792

/var/lib/openqa in QA-Power8-5-kvm reports Structure needs cleaning (auto_review:"(?s)Running on QA-Power8-5-kvm.*cannot remove directory for testresults":retry)

2020-10-29 15:04 - acarvajal

Status:	Closed	Start date:	2020-10-29
Priority:	Normal	Due date:	2020-11-06
Assignee:	acarvajal	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	future		

Description

Ran into several incomplete jobs in ppc64le, all of them in worker 1 from QA-Power8-5-kvm.

For example:

<https://openqa.suse.de/tests/4915970>
<https://openqa.suse.de/tests/4901191>
<https://openqa.suse.de/tests/4915970>
<https://openqa.suse.de/tests/4901090>

Cause of the failures seems to be that during job setup, worker encounters an existing testresults directory from a previous job.

Jobs also attempt to do a clean up of the testresults directory, but fail with error:

```
cannot remove directory for testresults: Directory not empty at /usr/lib/os-autoinst/bmwqemu.pm line 139.  
Can't mkdir('testresults'): File exists at /usr/lib/os-autoinst/bmwqemu.pm line 140
```

A look at the worker itself, reveals several files in /var/lib/openqa/pool/1/testresults with dirty structure:

```
QA-Power8-5-kvm:/var/lib/openqa/pool/1/testresults # ls -l  
ls: cannot access 'check_upgraded_service-230.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-527.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-585.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-612.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-299.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-664.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-212.png': Structure needs cleaning  
ls: cannot access 'check_upgraded_service-598.png': Structure needs cleaning  
total 36  
drwxr-xr-x 2 _openqa-worker nogroup 36864 Oct 29 14:54 .thumbs  
-????????? ? ? ? ? check_upgraded_service-212.png  
-????????? ? ? ? ? check_upgraded_service-230.png  
-????????? ? ? ? ? check_upgraded_service-299.png  
-????????? ? ? ? ? check_upgraded_service-527.png  
-????????? ? ? ? ? check_upgraded_service-585.png  
-????????? ? ? ? ? check_upgraded_service-598.png  
-????????? ? ? ? ? check_upgraded_service-612.png  
-????????? ? ? ? ? check_upgraded_service-664.png
```

History

#1 - 2020-10-29 15:04 - acarvajal

For the moment, stopped worker 1 in QA-Power8-5-kvm so it does not pick any jobs.

#2 - 2020-10-29 15:29 - acarvajal

Not sure if related: while checking status of the rest of the workers in QA-Power8-5-kvm.qa.suse.de (mainly ls and systemctl status in a shell), system became unresponsive.

It is currently not responding:

```
$ ping -c4 QA-Power8-5-kvm.qa.suse.de  
PING QA-Power8-5-kvm.qa.suse.de (10.162.6.202) 56(84) bytes of data.
```

--- QA-Power8-5-kvm.qa.suse.de ping statistics ---
4 packets transmitted, 0 received, 100% packet loss, time 3071ms

\$ ssh QA-Power8-5-kvm.qa.suse.de
ssh: connect to host qa-power8-5-kvm.qa.suse.de port 22: Connection timed out

#3 - 2020-10-29 16:09 - acarvajal

On connection via ipmitool to the console of QA-Power8-5-kvm saw the system booting.

Could see during boot errors mounting the NFS share /var/lib/openqa/share:

```
[FAILED] Failed to mount /var/lib/openqa/share.  
See 'systemctl status var-lib-openqa-share.mount' for details.  
[DEPEND] Dependency failed for Remote File Systems.
```

As well as FS errors from the local /var/lib/openqa:

```
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2280, block bitmap and bg descriptor inconsistent: 28857 vs 28876 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2259, block bitmap and bg descriptor inconsistent: 29490 vs 29515 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2272, block bitmap and bg descriptor inconsistent: 25721 vs 25727 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2276, block bitmap and bg descriptor inconsistent: 24337 vs 24292 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2278, block bitmap and bg descriptor inconsistent: 24534 vs 24323 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2279, block bitmap and bg descriptor inconsistent: 21293 vs 21305 free clusters  
EXT4-fs error (device sdb1): mb_free_blocks:1471: group 5094, inode 18039090: block 166924300:freeing already freed block (bit 4108); block  
bitmap corrupt.  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 5094, block bitmap and bg descriptor inconsistent: 32254 vs 33278 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 1488, block bitmap and bg descriptor inconsistent: 2094 vs 28078 free clusters  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 1487, block bitmap and bg descriptor inconsistent: 1534 vs 2958 free clusters  
EXT4-fs error (device sdb1): mb_free_blocks:1471: group 5678, inode 18038963: block 186074124:freeing already freed block (bit 17420); block  
bitmap corrupt.  
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 5678, block bitmap and bg descriptor inconsistent: 32254 vs 33278 free clusters  
EXT4-fs error (device sdb1): ext4_lookup:1588: inode #18047293: comm worker: deleted inode referenced: 18050333  
EXT4-fs error (device sdb1): ext4_lookup:1588: inode #18047293: comm worker: deleted inode referenced: 18051055  
EXT4-fs error (device sdb1): ext4_lookup:1588: inode #18047293: comm worker: deleted inode referenced: 18051173  
EXT4-fs error (device sdb1): ext4_lookup:1588: inode #18047293: comm worker: deleted inode referenced: 18051228  
EXT4-fs error: 33 callbacks suppressed  
EXT4-fs error (device sdb1): mb_free_blocks:1471: group 1464, inode 18038839: block 47972909:freeing already freed block (bit 557); block bitmap  
corrupt.
```

Since system had just come back, I proceeded to stop all workers as well as the openqa-worker-cacheservice, umount /var/lib/openqa and run an e2fsck for /dev/sdb1.

e2fsck corrected many errors:

- Inodes with wrong i_blocks.
- Inodes with zero dtime.
- File entries in /pool/1/testresults/.thumbs with deleted/unused inodes
- Unattached inodes.
- Block bitmap differences.
- Free block count wrong for several groups.
- Inode bitmap differences.
- Free inode count wrong for several groups.

A second call to e2fsck /dev/sdb1 showed no further issues:

```
QA-Power8-5-kvm:~ # e2fsck /dev/sdb1  
e2fsck 1.43.8 (1-Jan-2018)  
/dev/sdb1: clean, 61495/61054976 files, 21442894/244190384 blocks
```

After this was done:

- 1) Both /var/lib/openqa and /var/lib/openqa/share were mounted.
- 2) Offending files in /var/lib/openqa/pool/1/testresults were successfully removed.
- 3) Cacheservice was started
- 4) Workers were started.

Currently all 8 workers have picked up jobs:

<https://openqa.suse.de/tests/4915852#>
<https://openqa.suse.de/tests/4915847#>
<https://openqa.suse.de/tests/4915855#>
<https://openqa.suse.de/tests/4915854#>
<https://openqa.suse.de/tests/4915853#>
<https://openqa.suse.de/tests/4915851#>

<https://openqa.suse.de/tests/4915850#>
<https://openqa.suse.de/tests/4915848#>

I will monitor the jobs for a while.

#4 - 2020-10-29 17:52 - okurz

- Due date set to 2020-11-06
- Status changed from New to Feedback
- Assignee set to acarvajal
- Target version set to future

this is good progress. Actually I provided wrong information initially. I stated that the pool directory is initialized as a fresh filesystem on every boot. This is actually not the case for the powerVM machines (but all the others). So an explicit filesystem check and repair was certainly the right call. You mentioned that you will "monitor the jobs for a while" so I will assign the ticket to you and set the target version to "future" to show that it's actually not something that we in the SUSE QE Tools team need to actively followup with. I suggest the latest when the due date has passed without further problems you can resolve the ticket. Thanks so much already for helping us and helping you :)

#5 - 2020-10-30 12:09 - acarvajal

okurz wrote:

this is good progress. Actually I provided wrong information initially. I stated that the pool directory is initialized as a fresh filesystem on every boot. This is actually not the case for the powerVM machines (but all the others). So an explicit filesystem check and repair was certainly the right call. You mentioned that you will "monitor the jobs for a while" so I will assign the ticket to you and set the target version to "future" to show that it's actually not something that we in the SUSE QE Tools team need to actively followup with.

Perfect. I checked the first 8 jobs picked yesterday and they worked as expected, most with passing status and the ones that failed did so due to product issues and not due to openQA infrastructure.

I suggest the latest when the due date has passed without further problems you can resolve the ticket.

OK.

Thanks so much already for helping us and helping you :)

No problem.

#6 - 2020-10-30 13:41 - mkittler

- Subject changed from */var/lib/openqa in QA-Power8-5-kvm reports Structure needs cleaning* to */var/lib/openqa in QA-Power8-5-kvm reports Structure needs cleaning (auto_review:"(?s)Running on QA-Power8-5-kvm.*cannot remove directory for testresults")*

#7 - 2020-10-30 15:13 - okurz

- Subject changed from */var/lib/openqa in QA-Power8-5-kvm reports Structure needs cleaning (auto_review:"(?s)Running on QA-Power8-5-kvm.*cannot remove directory for testresults")* to */var/lib/openqa in QA-Power8-5-kvm reports Structure needs cleaning (auto_review:"(?s)Running on QA-Power8-5-kvm.*cannot remove directory for testresults":retry)*

I think "retry" is a good choice here because after retriggering any job will likely run elsewhere and not on the very same worker instance which is also fixed by now.

#8 - 2020-11-02 17:30 - acarvajal

Found today again QA-Power8-5-kvm unresponsive. Looking at the workers, it seems it didn't pick up any jobs after 31.10.2020 after noon:

<https://openqa.suse.de/admin/workers/1043>
<https://openqa.suse.de/admin/workers/1061>
<https://openqa.suse.de/admin/workers/1048>
<https://openqa.suse.de/admin/workers/1265>
<https://openqa.suse.de/admin/workers/1062>
<https://openqa.suse.de/admin/workers/1058>
<https://openqa.suse.de/admin/workers/1042>
<https://openqa.suse.de/admin/workers/1059>

I went to check on the monitoring site, and there is a gap between 30.10.2020 and 31.10.2020, followed by some activity on 31.10.2020 and then nothing after 31.10.2020:

<https://stats.openqa-monitor.qa.suse.de/d/WDQA-Power8-5-kvm/worker-dashboard-qa-power8-5-kvm?orgId=1&refresh=1m&from=now-7d&to=now>

I was able to power the system via IPMI, and on boot saw again problems with mounting /var/lib/openqa/share via NFS, as well as problems with sdb1:

```
QA-Power8-5-kvm login: EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2267, block bitmap and bg descriptor inconsistent: 30433 vs 30472 free clusters
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 4729, block bitmap and bg descriptor inconsistent: 22442 vs 22474 free clusters
BTRFS info (device sda2): relocating block group 588003672064 flags system|dup
BTRFS info (device sda2): relocating block group 587433246720 flags system|dup
contrack: generic helper won't handle protocol 47. Please consider loading the specific helper module.
BTRFS info (device sda2): relocating block group 589110968320 flags system|dup
```

```
QA-Power8-5-kvm login: BTRFS info (device sda2): relocating block group 589144522752 flags system|dup
BTRFS info (device sda2): relocating block group 589178077184 flags system|dup
EXT4-fs error (device sdb1): mb_free_blocks:1471: group 1100, inode 18038818: block 36046098:freeing already freed block (bit 1298); block bitmap corrupt.
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 1100, block bitmap and bg descriptor inconsistent: 1534 vs 1535 free clusters
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
EXT4-fs error (device sdb1): mb_free_blocks:1471: group 1842, inode 18038853: block 60367927:freeing already freed block (bit 9271); block bitmap corrupt.
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 1842, block bitmap and bg descriptor inconsistent: 20030 vs 20031 free clusters
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
EXT4-fs error (device sdb1) in ext4_do_update_inode:5272: error 27
BTRFS info (device sda2): relocating block group 589211631616 flags system|dup
BTRFS info (device sda2): relocating block group 587466801152 flags metadata|dup
BTRFS info (device sda2): found 1721 extents
EXT4-fs error (device sdb1): ext4_mb_generate_buddy:759: group 2023, block bitmap and bg descriptor inconsistent: 17608 vs 17646 free clusters
```

Not sure whether there is a problem with sdb1 that is making the system hang, or if the errors in sdb1 are product of the crash.

I repeated the procedure from 29.10.2020, i.e., stop services && umount /var/lib/openqa && e2fsck /dev/sdb1 && mount /var/lib/openqa && mount /var/lib/openqa/share && start services.

Workers are picking up jobs again, but definitely the system is not stable.

#9 - 2020-11-03 11:55 - okurz

mdoucha reported in chat that there are new incompletes due to corrupted cache sqlite db. Did:

```
systemctl stop salt-minion openqa-worker.target openqa-worker-cacheservice openqa-worker-cacheservice-minion.service telegraf.service && rm /var/lib/openqa/cache/* && systemctl start salt-minion openqa-worker.target openqa-worker-cacheservice openqa-worker-cacheservice-minion.service telegraf.service
```

Retriggered the last pipeline in <https://gitlab.suse.de/openqa/auto-review/-/pipelines> to cover any new unhandled incompletes for this worker.

#10 - 2020-12-03 12:53 - acarvajal

- % Done changed from 0 to 100

#11 - 2021-06-18 07:57 - acarvajal

- Status changed from Feedback to Closed