

openQA Infrastructure - action #80656

OSD deployment failed at 2020-12-02 because 'malbec.arch.suse.de' is down

2020-12-02 07:08 - Xiaojing_liu

Status:	Resolved	Start date:	2020-12-02
Priority:	Low	Due date:	
Assignee:	nicksinger	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
Observation			
OSD deployment failed, the error message is			
<pre>malbec.arch.suse.de: Minion did not return. [Not connected]</pre>			
Workaround			
Because malbec is down for a bit longer, so run salt-key -d malbec.arch.suse.de to remove malbec from salt-key on OSD. Then re-run the deployment.			
When malbec is back, we should add it into salt-key.			
Related issues:			
Related to openQA Infrastructure - action #81058: [tracker-ticket] Power mach...		Resolved	2020-12-15 2021-04-16

History

#1 - 2020-12-02 07:08 - Xiaojing_liu

OSD deployment has succeeded.

#2 - 2020-12-02 07:33 - Xiaojing_liu

- Status changed from Feedback to Workable
- Assignee deleted (Xiaojing_liu)

#3 - 2020-12-02 07:35 - nicksinger

- Assignee set to nicksinger

I'm currently in the process of recovering the machine. Afterwards I will re-add the salt-key on OSD

#4 - 2020-12-02 09:26 - nicksinger

I'm not sure how the machine booted previously. I *assume* we booted PXE and from there "timed out" into "boot from HDD". However, PXE reports an error now. I've created [RT-PPC #181643] and added osd-admin@suse.de as CC to address this issue. I will try to workaround the issue by manually kexec'ing the installed system as there is no dedicated bootloader entry for the installed system (just a installation entry).

#5 - 2020-12-02 09:39 - okurz

- Target version set to Ready

#6 - 2020-12-02 10:05 - nicksinger

Seems like a more severe issue. I can't find the systems boot disk at all:

```
/ # blkid
/dev/sdm1: UUID="6c7adfd9-8aa2-45e3-abf2-e6aff8ba8721"
/dev/sdf1: UUID="6c7adfd9-8aa2-45e3-abf2-e6aff8ba8721"
/dev/sdj1: UUID="e00ee584-c968-41c0-ab80-ad3ac3b68d97"
/dev/sdc1: UUID="e00ee584-c968-41c0-ab80-ad3ac3b68d97"
/ # mount /dev/sdm1 /mnt/tmp1/
```

```
/ # mount /dev/sdf1 /mnt/tmp2/
/ # mount /dev/sdj1 /mnt/tmp3/
/ # mount /dev/sdc1 /mnt/tmp4/
/ # find /mnt/tmp* -maxdepth 1
/mnt/tmp1
/mnt/tmp1/1
/mnt/tmp1/lost+found
/mnt/tmp2
/mnt/tmp2/1
/mnt/tmp2/lost+found
/mnt/tmp3
/mnt/tmp3/tmp
/mnt/tmp3/cache.sqlite-wal
/mnt/tmp3/cache.sqlite
/mnt/tmp3/cache.sqlite-shm
/mnt/tmp3/openqa.suse.de
/mnt/tmp3/lost+found
/mnt/tmp4
/mnt/tmp4/tmp
/mnt/tmp4/cache.sqlite-wal
/mnt/tmp4/cache.sqlite
/mnt/tmp4/cache.sqlite-shm
/mnt/tmp4/openqa.suse.de
/mnt/tmp4/lost+found
```

#7 - 2020-12-02 10:05 - nicksinger

- Status changed from Workable to In Progress

#8 - 2020-12-03 12:41 - nicksinger

```
[ 41.796572] Btrfs loaded
[ 41.797425] BTRFS: device fsid ae18adf5-d27e-4fa1-93a1-6ab55263c29d devid 1 transid 2545520 /dev/mapper/sdb
1
[ 41.798671] BTRFS info (device dm-5): disk space caching is enabled
[ 41.798675] BTRFS: has skinny extents
[ 42.663776] device-mapper: snapshots: Invalidating snapshot: Unable to allocate exception.
[ 42.665824] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 1, rd 0, flush 0, corrupt 0, gen 0
[ 42.667327] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 2, rd 0, flush 0, corrupt 0, gen 0
[ 42.667630] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 2, rd 1, flush 0, corrupt 0, gen 0
[ 42.667661] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 2, rd 2, flush 0, corrupt 0, gen 0
[ 42.667682] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 2, rd 3, flush 0, corrupt 0, gen 0
[ 42.667701] BTRFS error (device dm-5): bdev /dev/mapper/sdb1 errs: wr 2, rd 4, flush 0, corrupt 0, gen 0
[ 42.667714] BTRFS: Transaction aborted (error -5)
[ 42.667732] -----[ cut here ]-----
[ 42.667734] WARNING: at fs/btrfs/extent-tree.c:2930
[ 42.667735] Modules linked in: btrfs xor zlib_inflate lzo_compress lzo_decompress raid6_pq ext4 mbcache jbd
2 dm_snapshot dm_bufio dm_mod sd_mod sr_mod cdrom lpfc crc_t10dif crct10dif_generic crct10dif_common ipr
[ 42.667756] CPU: 146 PID: 2668 Comm: pb-discover Not tainted 4.4.92-openpower1 #1
[ 42.667758] task: c000001fc6241b40 ti: c000001fc6408000 task.ti: c000001fc6408000
[ 42.667760] NIP: d000000005acb344 LR: d000000005acb340 CTR: c0000000001e76a8
[ 42.667762] REGS: c000001fc640b240 TRAP: 0700 Not tainted (4.4.92-openpower1)
[ 42.667763] MSR: 900000000029033 <SF,HV,EE,ME,IR,DR,RI,LE> CR: 28042844 XER: 20000000
[ 42.667770] CFAR: c0000000005accd8 SOFTE: 1
[ 42.667770] GPR00: d000000005acb340 c000001fc640b4c0 d000000005b7ef58 0000000000000025
[ 42.667770] GPR04: 0000000000000001 00000000000004d4 0000000000000035 74726f6261206e6f
[ 42.667770] GPR08: 0000000000000007 0000000000000001 0000000000000007 000000000001ed60
[ 42.667770] GPR12: 0000000000002200 c00000000fe9b600 c000001fdd743100 c000001fdd743eb0
[ 42.667770] GPR16: c000001fdd7439c8 c000001fdd7439b0 0000000000001000 00000000026d770
[ 42.667770] GPR20: c000001fb0b401df c000001fb3df00e0 c000001fb3df0160 c000001fb3df0170
[ 42.667770] GPR24: c000001fb3c10088 0000000000000001 0000000000000000 0000000000000000
[ 42.667770] GPR28: c000001fb3df0000 c000001fb08a0800 c000001fb3c10000 ffffffffbbbbbb
[ 42.667808] NIP [d000000005acb344] btrfs_run_delayed_refs+0xdc/0x3a4 [btrfs]
[ 42.667818] LR [d000000005acb340] btrfs_run_delayed_refs+0xd8/0x3a4 [btrfs]
[ 42.667819] Call Trace:
[ 42.667828] [c000001fc640b4c0] [d000000005acb340] btrfs_run_delayed_refs+0xd8/0x3a4 [btrfs] (unreliable)
[ 42.667838] [c000001fc640b590] [d000000005accd1d0] btrfs_write_dirty_block_groups+0xdc/0x248 [btrfs]
[ 42.667847] [c000001fc640b650] [d000000005b52ccc] commit_cowonly_roots+0x208/0x2a8 [btrfs]
[ 42.667857] [c000001fc640b6e0] [d000000005adff80] btrfs_commit_transaction+0x5c0/0xaa4 [btrfs]
[ 42.667867] [c000001fc640b7b0] [d000000005ad9dd0] btrfs_commit_super+0xa0/0xac [btrfs]
[ 42.667877] [c000001fc640b7e0] [d000000005add474] open_ctree+0x1a7c/0x1dd8 [btrfs]
[ 42.667885] [c000001fc640b910] [d000000005ab2eb0] btrfs_remount+0xc08/0xee0 [btrfs]
[ 42.667889] [c000001fc640ba30] [c00000000011838c] mount_fs+0x94/0x174
[ 42.667892] [c000001fc640bac0] [c0000000001334a0] vfs_kern_mount+0x64/0x138
[ 42.667900] [c000001fc640bb10] [d000000005ab2950] btrfs_remount+0x6a8/0xee0 [btrfs]
```

```
[ 42.667903] [c000001fc640bc30] [c00000000011838c] mount_fs+0x94/0x174
[ 42.667905] [c000001fc640bcc0] [c0000000001334a0] vfs_kern_mount+0x64/0x138
[ 42.667907] [c000001fc640bd10] [c00000000013758c] do_mount+0xbcc/0xcfc
[ 42.667909] [c000001fc640bdd0] [c000000000137930] Sys_mount+0x90/0xc8
[ 42.667912] [c000001fc640be30] [c000000000009198] system_call+0x38/0xd0
[ 42.667914] Instruction dump:
[ 42.667916] 7d4048a8 7d474378 7ce049ad 40c2fff4 7c0004ac 7949f7e3 40e2001c 3c620000
[ 42.667920] e86384b8 7fe4fb78 4808d0b5 e8410018 <0fe00000> 3ca20000 e8a584c0 7fc3f378
[ 42.667924] ---[ end trace 91e6b5bb365bcd2 ]---
[ 42.667927] BTRFS: error (device dm-5) in btrfs_run_delayed_refs:2930: errno=-5 IO failure
[ 42.668000] BTRFS warning (device dm-5): Skipping commit of aborted transaction.
[ 42.668003] BTRFS: error (device dm-5) in cleanup_transaction:1746: errno=-5 IO failure
[ 42.668086] BTRFS error (device dm-5): cleaner transaction attach returned -30
[ 42.842396] BTRFS: open_ctree failed
[ 42.850783] device-mapper: snapshots: Snapshot is marked invalid.
[ 42.851754] Buffer I/O error on dev dm-8, logical block 1, async page read
[ 43.020417] device-mapper: snapshots: Snapshot is marked invalid.
[ 43.021087] EXT4-fs (dm-8): unable to read superblock
[ 43.219871] device-mapper: snapshots: Snapshot is marked invalid.
[ 43.220679] EXT4-fs (dm-8): unable to read superblock
```

it really looks bad on here. We might face some corruption...

#9 - 2020-12-11 12:35 - nicksinger

- Status changed from In Progress to Resolved

ok, I tried several times to re-assemble the disk setup to manually boot from petitboot. But I failed to make any progress here. I then realized that we have an "TW installer" boot entry in petitboot and booted that one to get a fully functional linux. With that I was able to successfully mount the root partition and chroot into it (see <https://wiki.gentoo.org/wiki/Chroot/en#Configuration> for preparation beforehand). Inside the chroot I did a zypper ref followed by a zypper dup which indeed updated quite some important tools:

```
2020-11-27 03:01:40|install|kernel-default|4.12.14-1p151.28.83.1|ppc64le||repo-update|b40de4df821b2c4597a46a456c41f6d0b3ff302f7a822288ec26c5fba04d1e2a|
2020-11-27 03:01:40|install|krb5|1.16.3-1p151.2.15.1|ppc64le||repo-update|a2de502a73819cef74d739690613b7df6fc4123ae771b729e6179c9d222fbcee|
2020-11-27 03:01:41|install|libcares2|1.17.0-1p151.3.6.1|ppc64le||repo-update|468b229c8df5314857ada47a6442812eaffabdd4362a590ad062ac46aca01863|
2020-11-27 03:01:41|install|libisc1606|9.16.6-1p151.11.15.1|ppc64le||repo-update|f732d30bfb95b4f808d69a1557761a4c8fc3134b452ccb1484cdb9bbb254e7f|
2020-11-27 03:01:41|install|libsystemd0|234-1p151.26.31.1|ppc64le||repo-update|850924f5a7a55c13086fbb94313765e81c99c908add1bf0bf4883037f0c9abb3|
2020-11-27 03:01:41|install|libudev1|234-1p151.26.31.1|ppc64le||repo-update|edd9e91ee47bd9d5b9d879a31f866b06578eded356a474e8ec19f032077a0ed7|
2020-11-27 03:01:41|install|pam|1.3.0-1p151.8.12.1|ppc64le||repo-update|46a6e55aab338a6668b4fb8751ea6121e8a4da90e78d61f7de8606f05cf53cb4|
2020-11-27 03:01:42|install|python3-bind|9.16.6-1p151.11.15.1|noarch||repo-update|3d263b63f79cddf735df83220dfb4bb6627e340ee4c23396569f751d429cfd95|
2020-11-27 03:01:42|install|systemd-bash-completion|234-1p151.26.31.1|noarch||repo-update|1d4a7b4b0093628607fa3f52a48f902b9214195a468893edae7e764770817844|
2020-11-27 03:01:43|install|libisc1600|9.16.6-1p151.11.15.1|ppc64le||repo-update|31a36946f34b5d58ff14124e588651a97cb93a0c8d114ced78e2b62663dac453|
2020-11-27 03:01:43|install|libdnsl605|9.16.6-1p151.11.15.1|ppc64le||repo-update|165d9cc4bc2cd6211fa0a8f7060fa802440e224129df6fb7be6d2110ae063f11|
2020-11-27 03:01:43|install|librados2|14.2.13.450+g65ea1b614d-1p151.2.28.1|ppc64le||repo-update|65ae00c01e0e1bfac2344611cdc84127fd24ef78507816171857c919bd7b5a91|
2020-11-27 03:01:43|install|libdevmapper-1_03|1.02.149-1p151.4.21.1|ppc64le||repo-update|70e1645b31798bafbe596162f1e1bf17028c23dc2c7974320e0b210659c3ef4c|
2020-11-27 03:01:43|install|sudo|1.8.22-1p151.5.9.1|ppc64le||repo-update|830b988abe35ddc71c2ee0afd4ec4d502b40ba56f31ea812be874de49f6bb13|
2020-11-27 03:01:49|install|systemd|234-1p151.26.31.1|ppc64le||repo-update|796d2c1d6a0f1d8ea3f0a0b41d6524743eae8f89e88fee790d71ada3bed72f3ec|
2020-11-27 03:01:49|install|libns1604|9.16.6-1p151.11.15.1|ppc64le||repo-update|fba8b21a5ec342637efb06c5fe996a647d7444e7a1a4806cfe7ebc9614683558|
2020-11-27 03:01:50|install|libiscfg1600|9.16.6-1p151.11.15.1|ppc64le||repo-update|b5c87febfd341daecb132f376ae95a51de6079227e3103ca98824e5114928dd2|
2020-11-27 03:01:50|install|librbd1|14.2.13.450+g65ea1b614d-1p151.2.28.1|ppc64le||repo-update|d7e1697804fc736096887fa2813b87f8e7bebae570837b2d18243e0a2296e0ce|
2020-11-27 03:01:50|install|libdevmapper-event1_03|1.02.149-1p151.4.21.1|ppc64le||repo-update|32a21c14a6c7763a716960f39ffa7bf44da394a366f0f4fd208a04218d92a7c9|
2020-11-27 03:01:52|install|udev|234-1p151.26.31.1|ppc64le||repo-update|60a8b488bdb43f084d99cf09b1603cd67a81c31285b44e9ae89d60d4c341da64|
2020-11-27 03:01:53|install|libirs1601|9.16.6-1p151.11.15.1|ppc64le||repo-update|0cd9e4ff010e150c399095ead8ff9d6e5da80f805b16f9c4525226b9a80555e|
2020-11-27 03:01:53|install|libbind9-1600|9.16.6-1p151.11.15.1|ppc64le||repo-update|42b3768dc0ffd85fa50237a8e6
```

```
db64235d1e77f11e4b0528882f001acdbf4756|
2020-11-27 03:01:53|install|liblvm2cmd2_02|2.02.180-lp151.4.21.1|ppc64le||repo-update|f4fca3ddef7b5e362f61f9e5
d1d158cc23432321b27e8891f1dcabaf1ceeaddf|
2020-11-27 03:01:53|install|liblvm2app2_2|2.02.180-lp151.4.21.1|ppc64le||repo-update|8a03cca4a470eaec59afbc04c
64467771d858c9814b40ee9a23ef49e68c810f6|
2020-11-27 03:01:55|install|systemd-network|234-lp151.26.31.1|ppc64le||repo-update|5dbc133a6d2eb6c9385290b65f2
eba19212cb17218403c1a883430af33eee19d|
2020-11-27 03:01:56|install|bind-utils|9.16.6-lp151.11.15.1|ppc64le||repo-update|26090e976db367c5c8cf18b62da85
e4f8c49cc4ee5f6667d92edf9a0a638eeac|
2020-11-27 03:01:56|install|systemd-sysvinit|234-lp151.26.31.1|ppc64le||repo-update|72f09c3154c6b4c3e6dbf465d1
bf7549da69f9abe4c4db047b89249bdb133acfa|
2020-11-27 03:01:57|install|device-mapper|1.02.149-lp151.4.21.1|ppc64le||repo-update|0fb1460f2b4acda31166656fc
9c52ce2504b784bccc201f2a2a20b48448f0600|
2020-11-27 03:01:59|install|lvm2|2.02.180-lp151.4.21.1|ppc64le||repo-update|9c584255f129e36d0b53f213ffb3268018
1f5cbcbce16d321b1ebc0c0f134ba6|
2020-11-27 03:01:59|install|kpartx|0.7.9+195+suse.16740c5-lp151.2.12.1|ppc64le||repo-update|25522085f70175e305
07051fcf2c7cfc2f7f00f6adbc331a865c70467e72a6a4|
2020-11-27 03:02:01|install|multipath-tools|0.7.9+195+suse.16740c5-lp151.2.12.1|ppc64le||repo-update|0c054c2ca
0379304fd108a6f54c8b2988a2b20a31afc846b3d32f07a22411eb|
```

regeneration of grub (and its config files) failed due to the missing /.snapshot folder. Inside the chroot I was able to rectify this with mount -o subvol=@/.snapshots /dev/sdb1 /.snapshots/.

After all this was done I tried to reboot (had to workaround that too with: <https://www.linuxjournal.com/content/rebooting-magic-way>) and malbec came back perfectly fine!

I checked all units which came all up successfully, added it back into salt and ran a manual high-state. Another reboot to validate worked successfully.

Unfortunately I have no clue *why* the system broke. But it seems to be back now.

#10 - 2020-12-13 12:18 - okurz

- Status changed from Resolved to Feedback

you have not mentioned if you put back the salt key so I checked that but it seems to be there. I could ping -4 malbec.arch but not login over ssh. <https://stats.openqa-monitor.qa.suse.de/d/4KkGdvvZk/osd-status-overview?viewPanel=70&orgId=1> shows it offline and unfortunately also with the ipmi commands specified in <https://gitlab.suse.de/openqa/salt-pillars-openqa/-/blob/master/openqa/workerconf.sls> I have currently no luck. Could you please check if you can reach it and teach me how to reliably interact with that machine? Also I suggest we only resolve this ticket after we could verify in like 3 reboots that the machine comes up fine again.

#11 - 2020-12-15 03:32 - nicksinger

okurz wrote:

you have not mentioned if you put back the salt key so I checked that but it seems to be there. I could ping -4 malbec.arch but not login over ssh. <https://stats.openqa-monitor.qa.suse.de/d/4KkGdvvZk/osd-status-overview?viewPanel=70&orgId=1> shows it offline and unfortunately also with the ipmi commands specified in <https://gitlab.suse.de/openqa/salt-pillars-openqa/-/blob/master/openqa/workerconf.sls> I have currently no luck. Could you please check if you can reach it and teach me how to reliably interact with that machine? Also I suggest we only resolve this ticket after we could verify in like 3 reboots that the machine comes up fine again.

yeah, seems like we *still* face v6 issues but this time we receive the wrong lease for the host. Maybe this caused that you couldn't connect with ssh onto the host?

This is also the reason why it shows as offline in our monitoring as telegraf on OSD tries to reach v6 (as a (wrong) AAAA record is present). I created [RT-ADM #182595] AutoReply: Host malbec.arch.suse.de receives wrong dhcpv6 lease - please update to address this issue and add our current DUID (can be found in /var/lib/wicked/duid.xml) onto arch's dhcpd6.

I guess this was caused by me. So definitely a good catch, thanks!

The ipmi issue I could reproduce. I logged onto malbec to cold reset the BMC and now it seems to be reachable again:

```
selenium ~ » ipmitool -I lanplus -C 3 -H fspl-malbec.arch.suse.de -P admin chassis power status
Chassis Power is on
```

Taking <https://progress.opensuse.org/issues/81020#note-4> into account I'd postpone the multiple reboots for now. But I keep the ticket here open now anyway.

#12 - 2020-12-15 03:48 - nicksinger

- Related to action #81058: [tracker-ticket] Power machines can't find installed OS. Automatic reboots disabled for now added

#13 - 2020-12-15 07:37 - nicksinger

- Status changed from Feedback to Blocked

Blocked until we receive an update from Infra about the v6 lease

#14 - 2020-12-17 12:59 - nicksinger

- Status changed from Blocked to Resolved

We have the machine and its lease back. Therefore

<https://stats.openga-monitor.qa.suse.de/d/4KkGdvvZk/osd-status-overview?viewPanel=70&orgId=1> shows as green again. BMC is wonky but nothing I can do in that regard. For now it works.