

openQA Infrastructure - action #96719

recover imagetester with broken filesystem/hardware (was: automatic updates on imagetester don't work and it failed to come up after reboot)

2021-08-10 14:36 - osukup

Status:	Resolved	Start date:	2021-07-29
Priority:	Normal	Due date:	2021-10-22
Assignee:	nicksinger	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
During work on https://progress.opensuse.org/issues/96311 , we found imagetester wasn't updated for 2 months investigate why wasn't automatic transactional update working and update imagetester.			
now blocked by https://infra.nue.suse.com/SelfService/Display.html?id=194271 , because it didn't survive reboot and this host hasn't any remote management interface			

History

#1 - 2021-08-10 14:56 - cdywan

- Subject changed from *imagetester automatic updates don't work* to *automatic updates on imagetester don't work and it failed to come up after reboot*
- Due date set to 2021-08-12
- Status changed from *New* to *Blocked*
- Priority changed from *Normal* to *High*

I'm bumping prio since the machine is completely offline right now and not just outdated. I made this clear in the title also. And *Blocked* with a *due date* on Thursday on the off chance that we forget about it.

#2 - 2021-08-11 14:45 - cdywan

- Blocks action #96311: *qemu error message is still "debug", should be "warn" or more severe size:S added*

#3 - 2021-08-11 14:48 - cdywan

Seems like there's several redundant disk errors on the machine that look like this and it's currently stuck in grub:

```
BTRFS error (device) md0p1): Remounting read0write after error is not allowed
```

#4 - 2021-08-13 09:59 - cdywan

- Due date changed from 2021-08-12 to 2021-08-20

cdywan wrote:

Seems like there's several redundant disk errors on the machine that look like this and it's currently stuck in grub:

```
BTRFS error (device) md0p1): Remounting read0write after error is not allowed
```

[osukup](#) reached out to @mrueckert to get access to the console

#5 - 2021-08-16 13:43 - cdywan

- Blocks deleted (action #96311: *qemu error message is still "debug", should be "warn" or more severe size:S*)

#6 - 2021-08-17 08:04 - osukup

infra sked for new disk:

but more importantly, can you organise a new disk and sent it to NUE that i can physically replace it?

#7 - 2021-08-17 08:45 - okurz

well, I guess we can get replacement hardware. We could wait for nsinger or try to go ahead ourselves and for example ask runger for help. I suggest to wait for nsinger to return from vacation and order together with him, ship to Nbg office and let EngInfra install the replacement hardware.

#8 - 2021-08-24 08:53 - nicksinger

a new disk sounds feasible. Just wondering: did we make sure the disk is actually broken? Or is it "just" the filesystem on there? Given that @mrueckert was involved I could imagine he checked but just to be sure :)

#9 - 2021-08-24 11:13 - osukup

[nicksinger](#) -> marcus ruecket isnt involved, record in rackspace is obsolete. Ticket is handled by [maxmaher](#) (Maximilian Maher), please contact him

#10 - 2021-08-25 12:22 - nicksinger

[osukup](#) could you please share this ticket? The ticket number, something?

I talked with gschlotter today to get access to the ipmi interface. Unfortunately it is an infra-only subnet we can't access. I currently still don't know what hardware we have in there...

#11 - 2021-08-25 12:44 - osukup

you mean <https://infra.nue.suse.com/SelfService/Display.html?id=194271> from description

#12 - 2021-08-25 12:50 - nicksinger

- Assignee changed from osukup to nicksinger

#13 - 2021-08-25 13:13 - nicksinger

Right, overlooked it. I've updated the ticket and talked to Max. I need IPMI access to the machine to continue further:

As discussed in RC it would be nice if somebody could reconfigure the switch so that we have access to the IPMI interface. This way I could 1) figure out if the HDD is actually broken or just the FS and 2) what is currently build in and what we need to buy. Since Max will be on vacation next week and might be too busy this week with other tasks it would be nice if somebody else from infra could take the reconfiguration of the switch.

I set this to blocked until this happened.

#14 - 2021-08-25 17:06 - cdywan

- Due date changed from 2021-08-20 to 2021-09-10

Moving due date as per conversation in chat since we're waiting on other people and it's not considered super urgent.

#15 - 2021-09-14 08:52 - okurz

- Due date changed from 2021-09-10 to 2021-09-17

- Status changed from Blocked to Feedback

[nicksinger](#) the infra ticket was resolved on 2021-09-07, so did you check if you do have IPMI access or something?

#16 - 2021-09-15 08:53 - nicksinger

- Status changed from Feedback to Blocked

unfortunately I didn't receive notifications. The ticket was closed with "please open a jira SD ticket". Done so now:

<https://sd.suse.com/servicedesk/customer/portal/1/SD-60360> (can anybody see this besides me?)

Nothing else happend.

#17 - 2021-09-15 13:27 - cdywan

nicksinger wrote:

unfortunately I didn't receive notifications. The ticket was closed with "please open a jira SD ticket". Done so now: <https://sd.suse.com/servicedesk/customer/portal/1/SD-60360> (can anybody see this besides me?)

I can't. Did you CC or otherwise add the ml or individual team members?

#18 - 2021-09-20 09:28 - nicksinger

Unfortunately I can only "share" the tickets with real accounts and not e-mails (like MLs). I added you, oli and marius for now manually.

#19 - 2021-09-20 09:57 - cdywan

- Due date changed from 2021-09-17 to 2021-09-22

Thanks! I can see it now. Bumping the due date to Wednesday for now.

#20 - 2021-09-22 15:06 - okurz

- Project changed from openQA Project to openQA Infrastructure

- Subject changed from automatic updates on imagetester don't work and it failed to come up after reboot to recover imagetester with broken filesystem/hardware (was: automatic updates on imagetester don't work and it failed to come up after reboot)

- Due date changed from 2021-09-22 to 2021-10-01

- Category deleted (Concrete Bugs)

- Status changed from Blocked to Feedback

I could actually access the machine now over IPMI (see SD ticket). I created https://gitlab.suse.de/openqa/salt-pillars-openqa/-/merge_requests/358

The SD ticket is still open and I asked for a proper DNS entry for IPMI. However, with the current state it should be possible to proceed hence "unblocking".

#21 - 2021-09-24 13:37 - nicksinger

- Status changed from Feedback to Workable

we could assign a *.qa.suse.de domain if nothing happens in the ticket.

#22 - 2021-09-29 09:13 - cdywan

Discussed it briefly in the unblock. We might use the IP or *.qam domain. Most importantly Nick will try and see how to restore the machine, maybe an office visit on Thursday

#23 - 2021-09-30 07:19 - nicksinger

I tried to access the machine over IPMI but apparently the Console redirection is misconfigured in the BIOS. Access over IPMIViewer is also not possible. So I need to check the machine in person today. Hopefully I catch somebody from infra who can let me into srv1.

#24 - 2021-10-01 06:50 - okurz

- Due date changed from 2021-10-01 to 2021-10-05

To be checked on-site at next possibility

#25 - 2021-10-03 17:55 - okurz

Regarding a good name for the IPMI endpoint maxmaher created https://gitlab.suse.de/OPS-Service/salt/-/merge_requests/1960 , not merged yet. It might be good for us to remember that repo so that at any next time we can create MRs ourselves

#26 - 2021-10-05 09:02 - okurz

- Due date changed from 2021-10-05 to 2021-10-22

- Priority changed from High to Normal

[nicksinger](#) as discussed in daily as you could not find anyone in office to collaborate on please create EngInfra ticket stating what goal we want to reach, e.g. check with IPMI command that we can read what's going on in the serial terminal, and suggest to change UEFI settings or something "on-site"

setting much further due-date as we rely a lot on individuals present on-site in nbg office and imagetester turned out to be not that critical right now, especially as we found already workarounds how to run openQA workers in containers (for s390x but we could apply the same elsewhere when we need to). Please still act urgently on raising the request with EngInfra, then we wait

#27 - 2021-10-05 11:27 - nicksinger

- Status changed from Workable to In Progress

Alright, after talking to [mkittler](#) today after the javaws stuff I had another idea how to access the iKVM of that machine. I went to the web interface of the BMC at <http://10.160.65.195> and clicked on the "Remote Console Preview" image there. This offers you to download a launch.jnlp file. Executing it on the console unfortunately fails with:

```

nsinger@workstation ~/Downloads » LANG=C javaws launch.jnlp
selected jre: /etc/java-config-2/current-system-vm/jre/
Warning!, Fall back in resolve_jar to hardcoded paths:
no
selected jre: /etc/java-config-2/current-system-vm/jre/
Warning!, Fall back in resolve_jar to hardcoded paths:
no
You are trying to get resource https://10.160.65.195:443/iKVM.jar but it is not in cache and could not be down
loaded. Attempting to continue, but you may expect failure
You are trying to get resource https://10.160.65.195:443/liblinux_x86_64.jar but it is not in cache and could
not be downloaded. Attempting to continue, but you may expect failure
JAR https://10.160.65.195:443/iKVM.jar not found. Continuing.
JAR https://10.160.65.195:443/liblinux_x86_64.jar not found. Continuing.
JAR https://10.160.65.195:443/iKVM.jar not found. Continuing.
JAR https://10.160.65.195:443/liblinux_x86_64.jar not found. Continuing.
netx: Initialization Error: Could not initialize application. (Fatal: Initialization Error: Unknown Main-Class
. Could not determine the main class for this application.)
net.sourceforge.jnlp.LaunchException: Fatal: Initialization Error: Could not initialize application. The appli
cation has not been initialized, for more information execute javaws from the command line.
    at net.sourceforge.jnlp.Launcher.createApplication(Launcher.java:822)
    at net.sourceforge.jnlp.Launcher.launchApplication(Launcher.java:531)
    at net.sourceforge.jnlp.Launcher$TgThread.run(Launcher.java:945)
Caused by: net.sourceforge.jnlp.LaunchException: Fatal: Initialization Error: Unknown Main-Class. Could not de
termine the main class for this application.
    at net.sourceforge.jnlp.runtime.JNLPCClassLoader.initializeResources(JNLPCClassLoader.java:774)
    at net.sourceforge.jnlp.runtime.JNLPCClassLoader.<init>(JNLPCClassLoader.java:338)
    at net.sourceforge.jnlp.runtime.JNLPCClassLoader.createInstance(JNLPCClassLoader.java:421)
    at net.sourceforge.jnlp.runtime.JNLPCClassLoader.getInstance(JNLPCClassLoader.java:495)
    at net.sourceforge.jnlp.runtime.JNLPCClassLoader.getInstance(JNLPCClassLoader.java:468)
    at net.sourceforge.jnlp.Launcher.createApplication(Launcher.java:814)
    ... 2 more

```

Following the leads in the error message I watched at an strace of the same "Remote console" but from openqaworker8. I grabbed its liblinux_x86_64.jar and iKVM.jar in the hopes that it might somehow work. It didn't. But the local webserver showed me what files imagetesters BMC actually requested:

```

2021-10-05 13:04:33.991 [INFO ] [::ffff:127.0.0.1]:35580 - HEAD /liblinux_x86_64__V1.0.3.jar.pack.gz (local: .
/liblinux_x86_64__V1.0.3.jar.pack.gz)
2021-10-05 13:04:33.991 [INFO ] [::ffff:127.0.0.1]:35582 - HEAD /iKVM__V1.69.13.0x0.jar.pack.gz (local: ./iKVM
__V1.69.13.0x0.jar.pack.gz)

```

With this information I was finally able to request the original files of from imagetester and put them into a temporary directory:

```

curl -k https://10.160.65.195:443/liblinux_x86_64__V1.0.3.jar.pack.gz > /tmp/ikvm/liblinux_x86_64__V1.0.3.jar.
pack.gz
curl -k https://10.160.65.195:443/iKVM__V1.69.13.0x0.jar.pack.gz > /tmp/ikvm/iKVM__V1.69.13.0x0.jar.pack.gz

```

I then started a webserver serving these files and modified the launch.jnlps first line from

```
<jnlp spec="1.0+" codebase="https://10.160.65.195:443/">
```

to:

```
<jnlp spec="1.0+" codebase="http://127.0.0.1:8888/">
```

Now I can start this modified jnlp-file with javaws launch.jnlp and it opens up a java application where I'm able to see the graphical output of imagetester. I can now start the investigation and we're not blocked by infra anymore.

#28 - 2021-10-05 13:28 - nicksinger

- Status changed from In Progress to Feedback

I booted a systemrescuecd and scrubbed the btrfs filesystem on the disk: no issues reported. Smart values look fine but error reporting doesn't seem to be supported so I couldn't check if there was any error reported in the past.

I then chrooted into the system (following <https://wiki.gentoo.org/wiki/Chroot/en#Configuration>):

```

chroot /mnt/mychroot
mount -a
transactional-update shell
zypper ref
zypper up

```

but it failed with:

```
Download (curl) error for 'http://download.opensuse.org/repositories/devel:/openQA/openSUSE_Leap_15.2/x86_64/0
```

```
s-autoinst-4.6.1633339717.7d37d2ac-lp152.875.1.x86_64.rpm':
Error code: Curl error 60
Error message: SSL certificate problem: self signed certificate in certificate chain
```

(all previous 227 packages where fine and no such error came up). I then tried to reboot the machine into the live-system again and it came back up. Not sure if the scrub fixed it or if transactional-upgrade did some magic to make it work again. On the live system I gave it another try with transactional-update up which successfully updated all installed packages. I then did another reboot and the machine came up successfully without any failing services in systemd. transactional-update.timer is also running too so lets see if the machine now works again.

#29 - 2021-10-05 13:33 - nicksinger

unfortunately I never got SOL to work. I checked the BIOS and everything looked fine regarding console redirection. I also tried to change values from (COM3* to COM2 and COM1) and several "IRQ" serial configurations. According to the docs of the BIOS the * at "COM3*" means this should be the console for SOL. I even got key presses redirected to the bios over ipmitool but nothing comes back in my local terminal. Also echo'ing stuff inside linux to the according /dev/ttyS* devices yields absolutely no output on my local ipmitool session. Given that we never even had BMC access before, I consider this "good enough" as I found a workaround for accessing the screen over this quirky [java "hack" mentioned above](#).

#30 - 2021-10-06 09:35 - cdywan

Please reboot it once more to ensure reboot stability, and then we can assume it "works".

#31 - 2021-10-06 09:41 - okurz

Nicely done. Impressive story and well done. I mentioned the note also in https://progress.opensuse.org/projects/openqav3/wiki/Wiki/diff?utf8=%E2%9C%93&version=128&version_from=127&commit=View+differences . Could you add a reference to our salt pillars file which also references o3 hosts?

#32 - 2021-10-06 10:42 - nicksinger

- Status changed from Feedback to Resolved

Did a transactional-update dup reboot, updates where successful and a reboot got scheduled:

```
imagetester:~ # rebootmgrctl is-active
RebootMgr is active
imagetester:~ # rebootmgrctl get-strategy
Reboot strategy: best-effort
imagetester:~ # rebootmgrctl get-window
Maintenance window is set to *--* 03:30:00, lasting 01h30m.
imagetester:~ # rebootmgrctl status
Status: Reboot requested, waiting for maintenance window
```

I canceled that reboot and triggered an immediate one: rebootmgrctl reboot now. After a few seconds the machine was back up and running.

#33 - 2021-10-06 11:32 - nicksinger

okurz wrote:

Nicely done. Impressive story and well done. I mentioned the note also in https://progress.opensuse.org/projects/openqav3/wiki/Wiki/diff?utf8=%E2%9C%93&version=128&version_from=127&commit=View+differences . Could you add a reference to our salt pillars file which also references o3 hosts?

Thanks! I made a more extensive wiki entry containing these information and referenced it with https://gitlab.suse.de/openqa/salt-pillars-openqa/-/merge_requests/359