

openQA Project - action #78390

coordination # 80142 (Blocked): [saga][epic] Scale out: Redundant/load-balancing deployments of openQA, easy containers, containers on kubernetes

Worker is stuck in "broken" state due to unavailable cache service (was: and even continuously fails to (re)connect to some configured web UIs)

2020-11-20 05:56 - okurz

Status:	Resolved	Start date:	2021-01-18
Priority:	Normal	Due date:	
Assignee:	mkittler	% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	Ready		
Difficulty:			

Description

Observation

On 2020-11-20 all worker instances on grenache-1 show up as "broken" and checked on grenache-1 with systemctl status openqa-worker@39 which says:

```
Nov 20 06:40:46 grenache-1 worker[6888]: [info] [pid:6888] Registering with openQA baremetal-support.qa.suse.de
Nov 20 06:40:46 grenache-1 worker[6888]: [warn] [pid:6888] Failed to register at baremetal-support.qa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Nov 20 06:40:56 grenache-1 worker[6888]: [info] [pid:6888] Registering with openQA baremetal-support.qa.suse.de
Nov 20 06:40:56 grenache-1 worker[6888]: [warn] [pid:6888] Failed to register at baremetal-support.qa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
```

so it *seems* like baremetal-support.qa.suse.de is not reachable. For more observations see the different comments.

Steps to reproduce

Likely can be reproduced by configuring a worker to connect to two web UI hosts where one can not be reached at all (e.g. valid DNS entry but host not up)

Acceptance criteria

- ~~AC1: The configured and reachable webUI hosts are reached while the down host is ignored as long as it is down~~
- ~~AC2: webUI does not show up as "broken" on a reachable webUI~~
- ~~AC3: worker still retries for multiple minutes when a webUI is temporarily down, e.g. during reboot~~

The fact that there are multiple web UIs involved is not really the issue here. See further comments.

Suggestions

Maybe when we extended the waiting periods while a webUI is down we introduced this regression

- DONE (at least regarding the initial suspicion): check history of our changes
- DONE (at least regarding the initial suspicion): crosscheck and extend tests
- fix behavior
- confirm working in OSD infrastructure
- DONE: remove workarounds, e.g. ensure baremetal-support.qa.suse.de is added back to salt-pillars-openqa for grenache-1

Workaround

Restart worker services

Related issues:

Related to openQA Infrastructure - action #80768: All workers in grenache-1 a...	Resolved	2020-12-07
Related to openQA Infrastructure - action #75238: Upgrade osd workers and ope...	Resolved	
Related to openQA Infrastructure - action #81046: openqaworker-arm-2.suse.de ...	Resolved	2020-12-15
Related to openQA Infrastructure - action #81210: workers in grenache-1 are b...	Resolved	2020-12-21
Copied from openQA Infrastructure - action #78218: [openQA][worker] Almost al...	Resolved	2020-11-19

History**#1 - 2020-11-20 05:56 - okurz**

- Copied from action #78218: [openQA][worker] Almost all openQA workers become offline added

#2 - 2020-11-25 12:08 - okurz

- Estimated time set to 80142.00 h

#3 - 2020-11-25 12:10 - okurz

- Estimated time deleted (80142.00 h)

#4 - 2020-11-25 12:16 - okurz

- Parent task set to #80142

#5 - 2020-12-02 12:53 - mkittler

- Assignee set to mkittler

#6 - 2020-12-02 15:18 - mkittler

Likely can be reproduced by configuring a worker to connect to two web UI hosts where one can not be reached at all (e.g. valid DNS entry but host not up)

I can not reproduce the problem that way. All the acceptance criteria are met - also when some of the configured web UI hosts are not reachable. Of course error messages for the failed registration attempts are logged continuously (also during the execution of a job for another web UI) but I don't consider that a problem.

Unfortunately the journal from the time the issue occurred is not present anymore. The ticket description also didn't state why the worker was broken at that time. However, the worker always adds an error message for its brokenness (e.g. cache service is down) and that would be very helpful to know. I also don't know any way the worker would declare itself broken because it can not connect to a web UI host. So I suppose the worker being broken and the failing attempts to register to another web UI were unrelated. We recently had problems with the SQLite database which could have been the cause of the broken state.

confirm working in OSD infrastructure
remove workarounds, e.g. ensure baremetal-support.qa.suse.de is added back to salt-pillars-openqa for grenache-1

At this point baremetal-support.qa.suse.de is back and grenache-1 registers correctly with it and OSD.

crosscheck and extend tests

We already have a very basic test for a configuration with multiple hosts but nothing where it will actually connect to any of the configured hosts.

Considering that it actually seems to work it looks more like the worker was broken for another reason - especially because the explicit "broken" state can so far only be reached if the cache service is unavailable, when there's a leftover QEMU process or when the pool directory is locked. There's really no code which would lead it to go into the "broken" state for anything else so far. So all I can do is extending the unit tests.

#7 - 2020-12-02 15:37 - mkittler

- Status changed from Workable to In Progress

PR for extending the unit tests (without much extra code): <https://github.com/os-autoinst/openQA/pull/3608>

#8 - 2020-12-03 04:07 - openqa_review

- Due date set to 2020-12-17

Setting due date based on mean cycle time of SUSE QE Tools

#9 - 2020-12-03 12:01 - okurz

PR with tests are merged. I read [#78390#note-6](#) . If you can't reproduce the issue easily maybe try to reproduce with the instances involved originally, e.g. coordinate with admins of baremetal-support.qa a controlled shutdown and check how the osd workers behave.

#10 - 2020-12-03 16:42 - mkittler

I'm wondering whether that's really worth the effort considering the broken state had very likely nothing to do with the failed connection attempts. There's simply no code which would bring connection failures and the "brokenness detection" together.

By the way, AC2 makes no sense if you read it on its own because it would mean removing the "broken" state completely. However, I assume that's not what we want.

And about AC3: So far the worker does not retry multiple minutes but actually endlessly. The only exception are authorization errors. But I guess that still counts as satisfying the AC.

#11 - 2020-12-07 11:08 - mkittler

Looks like the saw the problem on grenache-1 now again anyways. It showed up as broken on OSD and couldn't connect to baremetal-support.qa.suse.de at all.

The worker log looked like this:

```
Dez 07 11:44:56 grenache-1 worker[6746]: [error] [pid:6746] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported
Dez 07 11:44:56 grenache-1 worker[6746]: [info] [pid:6746] Registering with openQA baremetal-support.qa.suse.de
Dez 07 11:44:56 grenache-1 worker[6746]: [warn] [pid:6746] Failed to register at baremetal-support.qa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
...
Dez 07 11:51:31 grenache-1 worker[6746]: [error] [pid:6746] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported
Dez 07 11:51:31 grenache-1 worker[6746]: [info] [pid:6746] Registering with openQA baremetal-support.qa.suse.de
Dez 07 11:51:31 grenache-1 worker[6746]: [warn] [pid:6746] Failed to register at baremetal-support.qa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Dez 07 11:51:41 grenache-1 worker[6746]: [info] [pid:6746] Registering with openQA baremetal-support.qa.suse.de
Dez 07 11:51:41 grenache-1 worker[6746]: [warn] [pid:6746] Failed to register at baremetal-support.qa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
...
```

and the reason for the brokenness shown on OSD was Cache service info error: Can't connect: Address family for hostname not supported.

So as previously figured the worker did not show up as broken on OSD because it couldn't connect to the 2nd web UI host but because it couldn't connect to the cache service. After restarting the worker services both connections (cache service and baremetal-support.qa.suse.de) worked again. Not sure why that was necessary and the further connection attempts (for both connections) didn't work. In both cases the Mojolicious framework (and its underlying Perl modules) are used to connect. These apparently can get stuck in an error state. For now I've just restarted all worker slots again.

#12 - 2020-12-07 11:18 - kraih

The error Address family for hostname not supported comes from IO::Socket::IP. The fact that we keep getting the same error over and over probably means that it gets stuck somewhere.

#13 - 2020-12-07 11:22 - mkittler

By the way, it was **only** required to restart the workers. I haven't restarted the cache service or the web UI on baremetal-support.qa.suse.de. And according to its log the cache service was listening on both, IPv4 and IPv6:

```
Dez 06 03:35:35 grenache-1 openqa-workercache-daemon[6673]: Web application available at http://127.0.0.1:9530
Dez 06 03:35:35 grenache-1 openqa-workercache-daemon[6673]: Web application available at http://[::]:9530
```

#14 - 2020-12-07 16:11 - mkittler

And IO::Socket::IP just calls the c'tor of its base class IO::Socket. I couldn't find anything related in its [documentation](#). I've also checked [Socket](#). The error itself comes from the OS anyways so maybe there's even some strange caching on system level going on?

#15 - 2020-12-09 15:01 - mkittler

- Related to action #80768: All workers in grenache-1 are broken at 2020-12-07 added

#16 - 2020-12-10 08:10 - okurz

today grenache-1 workers report as "Broken", e.g. <https://openqa.suse.de/admin/workers/1207> , systemctl status shows:

```
# systemctl status openqa-worker@10
● openqa-worker@10.service - openQA Worker #10
   Loaded: loaded (/usr/lib/systemd/system/openqa-worker@10.service; enabled; vendor preset: disabled)
   Drop-In: /etc/systemd/system/openqa-worker@10.service.d
            └─30-openvswitch.conf
   Active: active (running) since Wed 2020-12-09 22:41:30 CET; 10h ago
   Main PID: 7213 (worker)
     Tasks: 1
   CGroup: /openqa.slice/openqa-worker.slice/openqa-worker@10.service
            └─7213 /usr/bin/perl /usr/share/openqa/script/worker --instance 10

Dec 10 09:06:38 grenache-1 worker[7213]: [info] [pid:7213] Registering with openQA baremetal-support.qa.suse.de
Dec 10 09:06:38 grenache-1 worker[7213]: [warn] [pid:7213] Failed to register at baremetal-support.qa.suse.de
- connection error: Can't connect: Name or service not known - trying again in 10 seconds
Dec 10 09:06:48 grenache-1 worker[7213]: [info] [pid:7213] Registering with openQA baremetal-support.qa.suse.de
Dec 10 09:06:48 grenache-1 worker[7213]: [warn] [pid:7213] Failed to register at baremetal-support.qa.suse.de
- connection error: Can't connect: Name or service not known - trying again in 10 seconds
Dec 10 09:06:58 grenache-1 worker[7213]: [info] [pid:7213] Registering with openQA baremetal-support.qa.suse.de
Dec 10 09:06:58 grenache-1 worker[7213]: [warn] [pid:7213] Failed to register at baremetal-support.qa.suse.de
- connection error: Can't connect: Name or service not known - trying again in 10 seconds
Dec 10 09:07:08 grenache-1 worker[7213]: [info] [pid:7213] Registering with openQA baremetal-support.qa.suse.de
Dec 10 09:07:08 grenache-1 worker[7213]: [warn] [pid:7213] Failed to register at baremetal-support.qa.suse.de
- connection error: Can't connect: Name or service not known - trying again in 10 seconds
Dec 10 09:07:18 grenache-1 worker[7213]: [info] [pid:7213] Registering with openQA baremetal-support.qa.suse.de
Dec 10 09:07:18 grenache-1 worker[7213]: [warn] [pid:7213] Failed to register at baremetal-support.qa.suse.de
- connection error: Can't connect: Name or service not known - trying again in 10 seconds
```

even though the error message might just be a symptom, we can at least try to improve that as well.

#17 - 2020-12-10 09:24 - cdywan

- Related to action #75238: Upgrade osd workers and openqa-monitor to openSUSE Leap 15.2 added

#18 - 2020-12-10 09:34 - mkittler

[okurz](#) But the "broken" state of the worker on OSD is *not* a symptom of the failed registration attempts to the 2nd web UI. The "broken" state of the worker on OSD is caused by the worker being unable to connect to the cache service. And that, in turn is caused by the same underlying connection problem which also let's the registration attempts fail. So no, I don't think the error message can be improved here. And if you would have looked a little bit further into the log you would have also spotted lines like Dec 10 09:38:48 grenache-1 worker[7238]: [error] [pid:7238] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported. These lines are the cause for the broken state on OSD.

#19 - 2020-12-10 09:50 - mkittler

grenache-1 has been restarted before I could restart the worker services. However, after the restart the problem persisted. So this problem maybe happens especially after a reboot. Possibly the service is started before the network is ready. Only after restarting the worker services via systemctl restart openqa-worker@{1..40} they behave as normal again.

#20 - 2020-12-10 09:54 - nicksinger

mkittler wrote:

grenache-1 has been restarted before I could restart the worker services. However, after the restart the problem persisted. So this problem maybe happens especially after a reboot. Possibly the service is started before the network is ready. Only after restarting the worker services via systemctl restart openqa-worker@{1..40} they behave as normal again.

kind of. After the report today that grenache-1 workers where offline again I assigned an v6 address to baremetal-support.qa.suse.de and added an according AAAA record to qanet's DNS server.

I then was fighting with the nsd service which cached the "NXDOMAIN" result for this webui. After restarting nsd the workers seem to be connected. I'm not sure if our actions crossed and your restarting of the service made them work again but I think it was just done automatically

#21 - 2020-12-10 12:34 - cdywan

So I also ran sudo systemctl restart openqa-worker@{1..10} on openqaworker10 and it seems to have brought those back *online* - but that's probably going to break after the next reboot ☐☐

#22 - 2020-12-10 16:04 - mkittler

What was wrong with openqaworker10? Did it show the same log/broken messages?

#23 - 2020-12-10 17:22 - cdywan

mkittler wrote:

What was wrong with openqaworker10? Did it show the same log/broken messages?

Yeah, it was showing up as *broken* in the web UI and had the network errors mentioned in this ticket. I saw this after rebooting it with 15.2 wondering why it wouldn't pick up jobs.

#24 - 2020-12-11 09:50 - mkittler

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

I've just checked myself. So yes, it is basically the same error. Since there's no 2nd web UI host configured on openqaworker10 we see of course only the errors concerning the cache service connection but it is the same underlying issue.

```
Dez 09 22:04:07 openqaworker10 worker[4883]: [error] [pid:4883] Worker cache not available: Cache service info
error: Can't connect: Address family for hostname not supported
Dez 09 22:04:07 openqaworker10 worker[4883]: [info] [pid:4883] Registered and connected via websockets with op
enQA host openqa.suse.de and worker ID 1314
Dez 09 22:04:07 openqaworker10 worker[4860]: [error] [pid:4860] Worker cache not available: Cache service info
error: Can't connect: Address family for hostname not supported
Dez 09 22:04:07 openqaworker10 worker[4860]: [info] [pid:4860] Registered and connected via websockets with op
enQA host openqa.suse.de and worker ID 1316
Dez 09 22:04:07 openqaworker10 worker[4856]: [error] [pid:4856] Worker cache not available: Cache service info
error: Can't connect: Address family for hostname not supported
Dez 09 22:04:07 openqaworker10 worker[4856]: [info] [pid:4856] Registered and connected via websockets with op
enQA host openqa.suse.de and worker ID 1312
Dez 09 22:04:07 openqaworker10 worker[4850]: [error] [pid:4850] Worker cache not available: Cache service info
error: Can't connect: Address family for hostname not supported
Dez 09 22:04:07 openqaworker10 worker[4850]: [info] [pid:4850] Registered and connected via websockets with op
enQA host openqa.suse.de and worker ID 1311
Dez 09 22:04:07 openqaworker10 worker[4895]: [error] [pid:4895] Worker cache not available: Cache service info
error: Can't connect: Address family for hostname not supported
Dez 09 22:04:07 openqaworker10 worker[4895]: [info] [pid:4895] Registered and connected via websockets with op
enQA host openqa.suse.de and worker ID 1313
```

Note that [nicksinger](#) started nscd at the same time when I restarted the workers as mentioned in #note-19. So it isn't clear to me whether that really fixed something but maybe it did so I keep the ticket in feedback.

#25 - 2020-12-11 10:00 - mkittler

- Project changed from *openQA Project* to *openQA Infrastructure*

- Subject changed from *worker is "broken" and not taking jobs if one of multiple configured webUI hosts is down to Worker is stuck in "broken" state and even continuously fails to (re)connect to some configured web UIs*

- Description updated

- Category deleted (*Concrete Bugs*)

#26 - 2020-12-13 12:23 - okurz

- Project changed from *openQA Infrastructure* to *openQA Project*

- Subject changed from *Worker is stuck in "broken" state and even continuously fails to (re)connect to some configured web UIs to Worker is stuck in "broken" state and even continuously fails to (re)connect to some configured web UIs*

- Category set to *Concrete Bugs*

please keep it in "openQA Project". We likely should improve the situation within openQA, at least with notes in the documentation about IPv4/IPv6 and/or the error messages

#27 - 2020-12-14 03:22 - Xiaojing_liu

This issue happens again in 2020-12-14 in some workers.

I tried to restart nscd firstly, but the workers' status is still broken, so I did `systemctl restart openqa-worker@[1..40]` in grenache-1.

I restarted the nscd at that time:

```
xiaojing@grenache-1:~> systemctl status nscd.service
● nscd.service - Name Service Cache Daemon
   Loaded: loaded (/usr/lib/systemd/system/nscd.service; enabled; vendor preset: enabled)
```

```
Active: active (running) since Mon 2020-12-14 04:03:29 CET; 11min ago
Process: 33430 ExecStop=/usr/sbin/nscd --shutdown (code=exited, status=0/SUCCESS)
Process: 33432 ExecStart=/usr/sbin/nscd (code=exited, status=0/SUCCESS)
```

And after 2020-12-14 04:03:29, the worker log showed [error] until restarting

```
Dec 14 03:54:57 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 03:59:52 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:04:47 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:09:43 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:12:06 grenache-1 worker[4656]: [info] [pid:4656] Quit due to signal TERM
Dec 14 04:12:06 grenache-1 worker[4656]: [debug] [pid:4656] Informing baremetal-support.qa.suse.de that we are going offline
Dec 14 04:12:06 grenache-1 worker[4656]: [debug] [pid:4656] Informing openqa.suse.de that we are going offline
Dec 14 04:12:06 grenache-1 systemd[1]: Stopping openQA Worker #10...
Dec 14 04:13:36 grenache-1 systemd[1]: openqa-worker@10.service: State 'stop-sigterm' timed out. Killing.
```

The same issue happened in arm-1, arm-2, openqaworker10, run `systemctl restart openqa-worker` on those hosts.

#28 - 2020-12-14 07:25 - nicksinger

Xiaojing_liu wrote:

This issue happens again in 2020-12-14 in some workers.

I tried to restart nscd firstly, but the workers' status is still broken, so I did `systemctl restart openqa-worker@{1..40}` in grenache-1.

I restarted the nscd at that time:

```
xiaojing@grenache-1:~> systemctl status nscd.service
● nscd.service - Name Service Cache Daemon
   Loaded: loaded (/usr/lib/systemd/system/nscd.service; enabled; vendor preset: enabled)
   Active: active (running) since Mon 2020-12-14 04:03:29 CET; 11min ago
   Process: 33430 ExecStop=/usr/sbin/nscd --shutdown (code=exited, status=0/SUCCESS)
   Process: 33432 ExecStart=/usr/sbin/nscd (code=exited, status=0/SUCCESS)
```

And after 2020-12-14 04:03:29, the worker log showed [error] until restarting

```
Dec 14 03:54:57 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 03:59:52 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:04:47 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:09:43 grenache-1 worker[4656]: [error] [pid:4656] Worker cache not available: Cache service info error: Can't connect: Address family for hostname >
Dec 14 04:12:06 grenache-1 worker[4656]: [info] [pid:4656] Quit due to signal TERM
Dec 14 04:12:06 grenache-1 worker[4656]: [debug] [pid:4656] Informing baremetal-support.qa.suse.de that we are going offline
Dec 14 04:12:06 grenache-1 worker[4656]: [debug] [pid:4656] Informing openqa.suse.de that we are going offline
Dec 14 04:12:06 grenache-1 systemd[1]: Stopping openQA Worker #10...
Dec 14 04:13:36 grenache-1 systemd[1]: openqa-worker@10.service: State 'stop-sigterm' timed out. Killing.
```

The same issue happened in arm-1, arm-2, openqaworker10, run `systemctl restart openqa-worker` on those hosts.

thanks for handling this.

Just to add some further information: I checked grenache-1 and it seems to have a completely working v4 *and* v6 stack. I can ping `::1` (localhost) as well as external resources (e.g. heise.de). [mkittler](#) Would it be possible to log the output of `ip a s` if we observe the error in the worker? Also: maybe an explicit connect to 127.0.0.1 instead of localhost could reduce potential issues? (not saying this is the right approach but maybe it yields some interesting results)

#29 - 2020-12-14 13:47 - mkittler

Would it be possible to log the output of `ip a s` if we observe the error in the worker?

I can do that before trying to apply any workarounds. However, it will likely only show that everything is configured correctly. But maybe at the time the worker services have been started something wasn't ready and for some reason the workers are stuck within that error state - even though the network should be fine at this point.

#30 - 2020-12-15 09:14 - cdywan

- Related to action #81046: `openqaworker-arm-2.suse.de` unreachable added

#31 - 2020-12-15 14:04 - okurz

- Subject changed from *Worker is stuck in "broken" state and even continuously fails to (re)connect to some configured web UIs* to *Worker is stuck in "broken" state due to unavailable cache service (was: and even continuously fails to (re)connect to some configured web UIs)*

Discussed with mkittler. mkittler explained me and convinced me that the original problem was really (only) that the worker cache service was not reachable. This was likely just due to a network misconfiguration problem. We have not found way to consistently reproduce this problem but we have found ideas for improving the log output. For this I opened <https://github.com/os-autoinst/openQA/issues/3645> and we coded together and came up with something that mkittler will create a PR for.

#32 - 2020-12-17 10:28 - mkittler

Now `openqaworker-arm-1` is affected (all active worker slots).

Unfortunately somebody else has just restarted all the worker slots before I could do any further investigation. It works again after that but I could not follow-up on [nicksinger](#) last comment. Note that `nscd` has *not* been restarted. So somehow I doubt that this service makes much of a difference here.

Here's just the log of one of the worker instances:

```
Dez 17 09:19:10 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:25:45 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:32:20 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:38:55 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:45:30 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:52:06 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 09:58:41 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 10:05:16 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 10:11:51 openqaworker-arm-1 worker[2553]: [error] [pid:2553] Worker cache not available: Cache service
info error: Can't connect: Address family for hostname not supported
Dez 17 10:14:24 openqaworker-arm-1 systemd[1]: Stopping openQA Worker #1...
Dez 17 10:14:24 openqaworker-arm-1 worker[2553]: [info] [pid:2553] Quit due to signal TERM
Dez 17 10:14:24 openqaworker-arm-1 worker[2553]: [debug] [pid:2553] Informing openqa.suse.de that we are going
offline
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: openqa-worker@1.service: State 'stop-sigterm' timed out. Killin
g.
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: openqa-worker@1.service: Killing process 2553 (worker) with sig
nal SIGKILL.
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: openqa-worker@1.service: Main process exited, code=killed, stat
us=9/KILL
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: Stopped openQA Worker #1.
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: openqa-worker@1.service: Unit entered failed state.
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: openqa-worker@1.service: Failed with result 'timeout'.
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: Starting openQA Worker #1...
Dez 17 10:15:54 openqaworker-arm-1 systemd[1]: Started openQA Worker #1.
```

Note that in this case there's no 2nd worker instances involved at all.

By the way, this does look like a general networking issue and is not necessarily IPv6 related. I assumed that the cache service client would try to use IPv6 but the code actually hard-codes an IPv4 address (`has host => sub { 'http://127.0.0.1' . service_port("cache_service") };` instead of `localhost`). Of course then it is weird that when it comes to connecting to multiple web UIs sometimes the connection to OSD works but not to other web UI hosts. I thought IPv4 vs. IPv6 might be responsible here.

#33 - 2020-12-17 10:42 - nicksinger

I need to stress again that restarting `nscd` was really only needed once to get the record for the kernel-webui recognized properly and flush the cached "NX AAAA record". This will not solve any issue in the future.

#34 - 2020-12-17 10:50 - nicksinger

I don't know how feasible this is but could we switch the communication between workers and the cacheservice to be done over unix domain sockets instead of local network communication? This would be a little bit more bulletproof in case of network misconfiguration.

#35 - 2020-12-18 10:18 - mkittler

I suppose it would be possible. I had to check whether there's something in Mojolicious/Minion for it.

Just for the record, it happened today again on OSD workers arm-1 and arm-2. On arm-1 restarting the worker services didn't even help unless also openqa-worker-cacheservice was restarted (no problems were visible on the journal of openqa-worker-cacheservice though).

#36 - 2020-12-21 06:15 - Xiaojing_liu

- Related to action #81210: workers in grenache-1 are broken added

#37 - 2020-12-22 11:01 - cdywan

- Due date deleted (2020-12-17)

#38 - 2021-01-14 11:51 - mkittler

Not sure whether there were any occurrences during my vacation but currently all slots on openqaworker-arm-1 are affected, e.g.:

```
-- Reboot --
Jan 14 11:12:08 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 14 11:12:59 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 14 11:13:00 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 14 11:13:01 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 14 11:13:03 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 14 11:13:42 openqaworker-arm-1 systemd[1]: Starting openQA Worker #1...
Jan 14 11:13:42 openqaworker-arm-1 systemd[1]: Started openQA Worker #1.
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [info] [pid:2549] worker 1:
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - config file: /etc/openqa/workers.ini
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - worker hostname: openqaworker-arm-1
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - isotovideo version: 20
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - websocket API version: 1
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - web UI hosts: openqa.suse.de
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - class: qemu_aarch64,qemu_aarch64_slow_worker,tap,openqaworker-arm-1
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - no cleanup: no
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: - pool directory: /var/lib/openqa/pool/1
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [error] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [info] [pid:2549] CACHE: caching is enabled, setting up /var/lib/openqa/cache/openqa.suse.de
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Project dir for host openqa.suse.de is /var/lib/openqa/share
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:13:52 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:02 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:02 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:12 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:12 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:22 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:22 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:32 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:32 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:42 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:42 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:14:52 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:14:52 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:15:02 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:15:02 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 14 11:15:12 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registering with openQA openqa.suse.de
Jan 14 11:15:12 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Establishing ws connection via ws://openqa.suse.de/api/v1/ws/476
Jan 14 11:15:13 openqaworker-arm-1 worker[2549]: [info] [pid:2549] Registered and connected via websockets wit
```



```

h openQA host openqa.suse.de and worker ID 476
Jan 14 11:15:13 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s
Jan 14 11:16:53 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s
Jan 14 11:20:13 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s
Jan 14 11:21:53 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s
Jan 14 11:23:33 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s
Jan 14 11:25:13 openqaworker-arm-1 worker[2549]: [warn] [pid:2549] Worker cache not available: Cache service info error: Can't connect: Address family for hostname not supported - checking again for web UI 'openqa.suse.de' in 100.00 s

```

Besides the actual problem, we have apparently a broken systemd override on that system.

The log clearly shows that when the worker service started the network wasn't up and both - the local connection to the cache service failed (with "Address family for hostname not supported") and the connection to openqa.suse.de (DNS resolution failed). The remote connection starts to work on its own again after several retries as expected. Only the local connection stays broken despite several retry attempts. At least we can observe that <https://github.com/os-autoinst/openQA/pull/3646> works.

And of course the cache service itself is accessible via IPv4 and IPv6 (the worker uses IPv4):

```

> curl "http://[::1]:9530/foo"
Not found
> curl "http://127.0.0.1:9530/foo"
Not found

```

I straced the worker process (there is only one process per slot; no forking happens in that state). During the "Cache service info error" nothing is happening. On a connection attempt, the following is logged:

```

> sudo strace -p 2549
...
getpid() = 2549
ppoll([{fd=3, events=POLLIN|POLLPRI|POLLOUT}], 1, {tv_sec=500, tv_nsec=178000000}, NULL, 0) = 1 ({fd=3, revents=POLLOUT}), left {tv_sec=500, tv_nsec=177996730}
write(3, "\201\376\0\244\0{7\212{YE\357a\10X\344\"A\25\335o\t\\\357r[T\353c\23R\252\"...", 172) = 172
getpid() = 2549
flock(1, LOCK_EX) = 0
write(1, "[warn] [pid:2549] Worker cache n\"...", 186) = 186
flock(1, LOCK_UN) = 0
ppoll([{fd=3, events=POLLIN|POLLPRI}], 1, {tv_sec=100, tv_nsec=0}, NULL, 0) = 1 ({fd=3, revents=POLLIN}), left {tv_sec=99, tv_nsec=742701000}
read(3, "\201 {\"population\":335,\"type\": \"info\"...", 131072) = 34
ppoll([{fd=3, events=POLLIN|POLLPRI}], 1, {tv_sec=99, tv_nsec=743000000}, NULL, 0

```

Only the 2nd and 3rd line are interesting and it repeats endlessly like that.

The descriptor 3 is apparently some socket:

```

sudo ls -l /proc/2549/fd
insgesamt 0
lr-x----- 1 _openqa-worker nogroup 64 14. Jan 11:13 0 -> /dev/null
lrwx----- 1 _openqa-worker nogroup 64 14. Jan 11:13 1 -> 'socket:[67656]'
lrwx----- 1 _openqa-worker nogroup 64 14. Jan 11:13 2 -> 'socket:[67656]'
lrwx----- 1 _openqa-worker nogroup 64 14. Jan 11:13 3 -> 'socket:[85384]'

```

But sudo lsof -i -a -p 2549 just hanngs and show nothing and it seems not possible to get further info:

```

> sudo head /proc/2549/fd/3
head: '/proc/2549/fd/3' kann nicht zum Lesen geöffnet werden: Kein passendes Gerät bzw. keine passende Adresse gefunden

```

I get the same result for the other FDs and also for the other worker slots.

#39 - 2021-01-14 12:40 - mkittler

By the way, the code which checks for the cache service is `$cache_service_client->info->availability_error` where `$cache_service_client` is always the same `OpenQA::CacheService::Client` instance which uses always the same `Mojo::UserAgent` instance under the hood (`$ix = $ua->$method(@args)`).

I've just restarted one worker slot and as usual that helps. In the good case the trace for requests to the cache service looks like this:

```
socket(AF_INET, SOCK_STREAM, IPPROTO_TCP) = 3
ioctl(3, TCGETS, 0xfffffe593f420) = -1 ENOTTY (Unpassender IOCTL (I/O-Control) für das Gerät)
lseek(3, 0, SEEK_CUR) = -1 ESPIPE (Nicht erlaubter Seek)
ioctl(3, TCGETS, 0xfffffe593f420) = -1 ENOTTY (Unpassender IOCTL (I/O-Control) für das Gerät)
lseek(3, 0, SEEK_CUR) = -1 ESPIPE (Nicht erlaubter Seek)
fcntl(3, F_SETFD, FD_CLOEXEC) = 0
fcntl(3, F_GETFL) = 0x2 (flags O_RDWR)
fcntl(3, F_SETFL, O_RDWR|O_NONBLOCK) = 0
connect(3, {sa_family=AF_INET, sin_port=htons(9530), sin_addr=inet_addr("127.0.0.1")}, 16) = -1 EINPROGRESS (Die Operation ist jetzt in Bearbeitung)
fcntl(3, F_GETFL) = 0x802 (flags O_RDWR|O_NONBLOCK)
ppoll([{fd=3, events=POLLOUT}], 1, {tv_sec=299, tv_nsec=999000000}, NULL, 0) = 1 ([{fd=3, revents=POLLOUT}], 1, left {tv_sec=299, tv_nsec=998996819})
getsockopt(3, SOL_SOCKET, SO_ERROR, [0], [256->4]) = 0
connect(3, {sa_family=AF_INET, sin_port=htons(9530), sin_addr=inet_addr("127.0.0.1")}, 16) = 0
getpeername(3, {sa_family=AF_INET, sin_port=htons(9530), sin_addr=inet_addr("127.0.0.1")}, [256->16]) = 0
setsockopt(3, SOL_TCP, TCP_NODELAY, [1], 4) = 0
getsockname(3, {sa_family=AF_INET, sin_port=htons(57106), sin_addr=inet_addr("127.0.0.1")}, [256->16]) = 0
getsockname(3, {sa_family=AF_INET, sin_port=htons(57106), sin_addr=inet_addr("127.0.0.1")}, [256->16]) = 0
getpeername(3, {sa_family=AF_INET, sin_port=htons(9530), sin_addr=inet_addr("127.0.0.1")}, [256->16]) = 0
ppoll([{fd=3, events=POLLIN|POLLPRI|POLLOUT}], 1, {tv_sec=10, tv_nsec=0}, NULL, 0) = 1 ([{fd=3, revents=POLLOUT}], 1, left {tv_sec=9, tv_nsec=999996890})
write(3, "GET /info HTTP/1.1\r\nAccept-Encod...", 118) = 118
ppoll([{fd=3, events=POLLIN|POLLPRI}], 1, {tv_sec=10, tv_nsec=0}, NULL, 0) = 1 ([{fd=3, revents=POLLIN}], 1, left {tv_sec=9, tv_nsec=996292269})
read(3, "HTTP/1.1 200 OK\r\nContent-Length:...", 131072) = 326
```

I'm wondering why the stuck workers don't even show a connect(3, {sa_family=AF_INET, sin_port=htons(9530), sin_addr=inet_addr("127.0.0.1")}, 16) like here. Maybe Perl or more likely getaddrinfo() remembers the last "Address family for hostname not supported" error and therefore it doesn't even try. (getaddrinfo() is not a syscall and therefore doesn't show up in strace.)

Likely we can not do anything about it except depending on network-online.target within the systemd service (to pull in wicked.service in our case). The worker currently only depends on network.target (which only pulls in openswitch.service).

#40 - 2021-01-14 15:03 - mkittler

If the problem is really getaddrinfo() being stuck in that error state there's likely nothing we can do to improve except trying to make sure we don't run into the error in the first place. Hence I created a PR to make the worker depend on network-online.target:

<https://github.com/os-autoinst/openQA/pull/3673>

#41 - 2021-01-15 12:24 - mkittler

We would not be the only ones running into that problem:

- <https://ask.puppet.com/question/11481/name-or-service-not-known-until-puppet-restart>
- <https://unix.stackexchange.com/questions/89699/process-calling-getaddrinfo-at-boot-gets-permanently-stuck-with-bad-etc-resol>
 - Apparently it helped to call `res_init()` which I can give a try as well.
 - It is suggested to use a local DNS resolver. We indeed don't do that already. However, I don't think it would help with the "127.0.0.1/"Address family for hostname not supported" case.
- <https://github.com/nning/transmission-rss/issues/42>
 - The resolution was to rely on network-online.target.

At least for the issue with the local connection (and that seems to be the only connection which still shows this issue) we could also avoid calling getaddrinfo() at all. This would require some changes in Perl libraries we're using. Some notes from [krah](#) about this:

```
fwiw. if you wanted to optimize getaddrinfo away in some cases, i believe this is the place you would have to do it
https://metacpan.org/source/PEVANS/IO-Socket-IP-0.41/lib/IO/Socket/IP.pm#L516
there is a possibility to optimize it on the mojo layer too, since IO::Socket::IP accepts a config option (PeerAddrInfo) with an already prepared
getaddrinfo data structure
we use that to integrate Net::DNS::Native (async resolver)
https://github.com/mojolicious/mojo/blob/296c11e4b2ee82b7609a7a0ec8d7696617bb9e69/lib/Mojo/IOLoop/Client.pm#L52
The fake version for platforms without getaddrinfo might be helpful too https://metacpan.org/release/Socket/source/Socket.pm#L970
guess for the simple 127.0.0.1 case you could just get away with pack_sockaddr_in
guess i could add an upstream feature to mojo to allow socket options to be passed with the transaction
$tx->socket_options({PeerAddrInfo => $already_prepared_data_structure});
that's quite a bit of work though (for upstream testing and docs), so we should be fairly certain we really really want this before i invest the time
```

#42 - 2021-01-18 09:17 - mkittler

The workers on openqaworker5 were broken this morning because of Address family for hostname not supported, I just restart all the worker services.

#43 - 2021-01-18 09:23 - mkittler

PR for a workaround from our side: <https://github.com/os-autoinst/openQA/pull/3676>

#44 - 2021-01-18 15:06 - okurz

I rebooted openqaworker6 and openqaworker10 today after a user report in <https://chat.suse.de/channel/testing?msg=udQguXCPNRcAABnBg>

#45 - 2021-01-19 13:07 - okurz

We can await the next osd deployment, then check for "broken" workers on <https://openqa.suse.de/admin/workers>

#46 - 2021-01-20 16:38 - mkittler

I've checked openqaworker-arm-1 and 2 and found no occurrences so far (for i in {1..10}; do sudo journalctl --no-pager --since '24 hours ago' -u openqa-worker@\$i | grep -i 'Address family for hostname'; done).

#47 - 2021-01-25 10:15 - mkittler

I haven't seen any broken workers anymore. Note that my greping on the journal was likely not very useful because the remaining errors might have a different message. The message Name or service not known is still being logged but the worker is not stuck in that error state so that should not be a problem. The new debug message Running patched getaddrinfo() for 127.0.0.1 introduced by my workaround is also visible in the log. It clutters the log a little bit so we might consider removing that message. That's how it looks like:

```
[...]
Jan 24 03:39:32 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:32 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:39:42 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:42 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:39:52 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:52 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:40:02 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:40:02 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:40:12 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:40:12 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Establishing ws connection via ws://openqa.suse.de/api/v1/ws/482
Jan 24 03:40:13 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:40:13 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registered and connected via websockets with openQA host openqa.suse.de and worker ID 482
Jan 24 03:41:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:43:01 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 24 03:43:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:45:13 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:46:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:48:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:50:13 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:51:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:53:34 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:55:14 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:56:54 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:58:34 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 04:00:14 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
```



```

Jan 24 12:22:22 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:24:02 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:25:42 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:27:22 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:29:02 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:30:42 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:32:22 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:34:02 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:35:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:37:23 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:39:03 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:40:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:42:23 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:44:03 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:45:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:47:23 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:49:03 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:50:44 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:52:24 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:54:04 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:55:44 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:57:24 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 12:59:04 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:00:44 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:02:24 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:04 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Accepting job 5334244 from openqa.suse.de.
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Setting job 5334244 from openqa.suse.de up
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Preparing Mojo::IOLoop::ReadWriteProcess::Session
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [info] [pid:2132] +++ setup notes +++
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Running on openqaworker-arm-1:7 (Linux 5.8.3-1.gbad027a-default #1 SMP Sat Aug 22 06:31:07 UTC 2020 (bad027a) aarch64)
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Job settings:
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132]
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: AGGREGATE_NEEDED=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: ARCH=aarch64
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: BACKEND=qemu
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: BOOT_HDD_IMAGE=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: BUILD=:17943:sudo
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: BUILD_SDK=GM
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: CACHEDIRECTORY=/var/lib/openqa/cache
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: COLLECT_COREDUMPS=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: DESKTOP=textmode
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: DISTRI=sle
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: FLAVOR=Server-DVD-Incidents-Install
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: HDDMODEL=virtio-blk-device
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: HDD_1=SLES-12-SP5-aarch64-Installtest.qcow2
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: INCIDENT_ID=17943

```

```

Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: INCIDENT_REPO=http://download.suse.de/ibs/SUSE:/Maintenance:/17943/SUSE_Updates_SLE-SDK_12-SP5_aarch64/,http://download.suse.de/ibs/SUSE:/Maintenance:/17943/SUSE_Updates_SLE-SERVER_12-SP5_aarch64/
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: INSTALLTEST=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: ISO=SLE-12-SP5-Server-DVD-aarch64-GM-DVD1.iso
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: JOBTOKEN=AzV8fhvPrcjGhUfN
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: LOG_LEVEL=debug
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: MACHINE=aarch64-virtio
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: MAX_JOB_TIME=9000
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: NAME=05334244-sle-12-SP5-Server-DVD-Incidents-Install-aarch64-Build:17943:sudo-qam-incidentinstall@aarch64-virtio
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: NOVIDEO=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: OPENQA_HOSTNAME=openqa.suse.de
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: OPENQA_URL=openqa.suse.de
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: OS_TEST_ISSUES=17943
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: PRJDIR=/var/lib/openqa/share
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMU=aarch64
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMUCPU=host
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMUCPUS=2
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMUMACHINE=virt,usb=off,gic-version=3,its=off
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMUPORT=20072
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: QEMU_DISABLE_SNAPSHOTS=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: REPO_ID=1611490232
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: RETRY_DELAY=5
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: RETRY_DELAY_IF_WEBUI_BUSY=60
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: SDK_TEST_ISSUES=17943
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: SERIALDEV=ttyAMA0
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: SHUTDOWN_NEEDS_AUTH=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: TEST=qam-incidentinstall
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: TEST_SUITE_NAME=qam-incidentinstall
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: UEFI=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: UEFI_PFLASH_CODE=/usr/share/qemu/aavmf-aarch64-code.bin
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: UEFI_PFLASH_VARS=SLES-12-SP5-aarch64-Installtest-uefi-vars.qcow2
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: VERSION=12-SP5
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: VGA=std
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: VIRTIO_CONSOLE=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: VNC=97
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: VNC_TYPING_LIMIT=10
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: WORKER_CLASS=qemu_aarch64,qemu_aarch64_slow_worker,tap,openqaworker-arm-1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: WORKER_HOSTNAME=10.160.0.245
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: WORKER_ID=482
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: WORKER_INSTANCE=7
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: ZYPPEL_ORPHANED_CHECK_ONLY=1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Found HDD_1, caching SLES-12-SP5-aarch64-Installtest.qcow2
Jan 24 13:04:28 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Downloading SLES-12-SP5-aarch64-Installtest.qcow2, request #13 sent to Cache Service
Jan 24 13:04:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Updating status so job 5334244 is not considered dead.
Jan 24 13:04:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5334244/status
Jan 24 13:04:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:38 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Updating status so job 5334244 is not considered dead.
Jan 24 13:04:38 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5334244/status
Jan 24 13:04:38 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Updating status so job 5334244 is not considered dead.
Jan 24 13:04:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5334244/status
Jan 24 13:04:43 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:48 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Updating status so job 5334244 is not considered dead.
Jan 24 13:04:48 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5334244/status
Jan 24 13:04:48 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1

```

```
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Updating status so job 5334244 is not considered dead.
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5334244/status
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Download of SLES-12-SP5-aarch64-Installtest.qcow2 processed:
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [info] [#13] Cache size of "/var/lib/openqa/cache" is 12GiB, with limit 50GiB
[...]
```

#48 - 2021-01-25 11:24 - cdywan

mkittler wrote:

```
[...]
Jan 24 03:39:32 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:32 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:39:42 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:42 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:39:52 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:39:52 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:40:02 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:40:02 openqaworker-arm-1 worker[2132]: [warn] [pid:2132] Failed to register at openqa.suse.de - connection error: Can't connect: Name or service not known - trying again in 10 seconds
Jan 24 03:40:12 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registering with openQA openqa.suse.de
Jan 24 03:40:12 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Establishing ws connection via ws://openqa.suse.de/api/v1/ws/482
Jan 24 03:40:13 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:40:13 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Registered and connected via websockets with openQA host openqa.suse.de and worker ID 482
Jan 24 03:41:53 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
Jan 24 03:43:01 openqaworker-arm-1 systemd[1]: /etc/systemd/system/openqa-worker@.service.d/30-openvswitch.conf:2: Unknown lvalue 'Environment' in section 'Unit'
Jan 24 03:43:33 openqaworker-arm-1 worker[2132]: [debug] [pid:2132] Running patched getaddrinfo() for 127.0.0.1
[...]
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [info] [pid:2132] Download of SLES-12-SP5-aarch64-Installtest.qcow2 processed:
Jan 24 13:04:53 openqaworker-arm-1 worker[2132]: [info] [#13] Cache size of "/var/lib/openqa/cache" is 12GiB, with limit 50GiB
[...]
```

Yeah, looks pretty good. You can see how's waiting on the connection. If this isn't debug logging, it probably should be, though.

Btw what's this Unknown lvalue 'Environment' in section 'Unit'? I guess that's unrelated, but since it is related to network setup I figured I'll ask anyway. I can't find a ticket about it.

#49 - 2021-01-28 12:44 - mkittler

Btw what's...

I've already created a SR for this and cleaned up the file at its old location. So it is taken care of.

#50 - 2021-02-01 11:29 - mkittler

PR for not relying on the monkey patch anymore with Mojolicious 8.72: <https://github.com/os-autoinst/openQA/pull/3704>

I would mark this ticket as resolved as soon as the PR has been merged. Of course the monkey patch is still present but we can remove it easily as soon as we no longer support Mojolicious 8.72.

#51 - 2021-02-01 11:41 - cdywan

mkittler wrote:

PR for not relying on the monkey patch anymore with Mojolicious 8.72: <https://github.com/os-autoinst/openQA/pull/3704>

I would mark this ticket as resolved as soon as the PR has been merged. Of course the monkey patch is still present but we can remove it easily as soon as we no longer support Mojolicious 8.72.

Sounds sensible to me ☐☐

#52 - 2021-02-02 20:06 - okurz

<https://github.com/os-autoinst/openQA/pull/3704> is merged. If we are lucky this will be automatically deployed to osd tomorrow morning. So I suggest you give the infrastructure a quick check before resolving.

#53 - 2021-02-04 15:22 - mkittler

I haven't seen any broken workers so far. I've also been checking the journal of our three OSD arm workers (where the problem occurred especially frequently) and couldn't find anything. The line "Running patched..." is no longer logged after the deployment yesterday which makes sense as the workers are already at Mojolicious 8.72. So my last PR could take effect and the previously added monkey patch is no longer used.

#54 - 2021-02-04 20:09 - okurz

- *Status changed from Feedback to Resolved*

- *Priority changed from Normal to Low*

Very good. So let's assume it's resolved then, right? I am very happy to see a solution which improved upstream as well. Good job everyone! :)

I have moved the sub-ticket up one level so that we can close this story as you continued here before closing the subtask.

#55 - 2021-02-04 20:09 - okurz

- *Priority changed from Low to Normal*