

## openQA Infrastructure - action #73633

### OSD partially unresponsive, triggering 500 responses, spotty response visible in monitoring panels but no alert triggered (yet)

2020-10-20 13:53 - okurz

<b>Status:</b>	Resolved	<b>Start date:</b>	2020-10-20
<b>Priority:</b>	Urgent	<b>Due date:</b>	2020-11-17
<b>Assignee:</b>	nicksinger	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Ready		
<b>Description</b>			
<b>Observation</b>			
<p><a href="https://monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&amp;from=1603190156643&amp;to=1603196975018">https://monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&amp;from=1603190156643&amp;to=1603196975018</a> shows that at around 2020-10-20 12:39 the HTTP response time from osd increased and users reported spotty connection and 500 responses "unresponsive" during that time, e.g. in <a href="https://chat.suse.de/channel/testing?msg=aix9KNXwkWowTd7FA">https://chat.suse.de/channel/testing?msg=aix9KNXwkWowTd7FA</a>. The spotty response is visible in our monitoring panels we no alert triggered so far in grafana because we do not want the unspecific "No Data" alerts.</p>			
<b>Cause, solution and test</b>			
<ul style="list-style-type: none"><li>• What caused this: <a href="https://progress.opensuse.org/issues/73633#note-17">https://progress.opensuse.org/issues/73633#note-17</a></li><li>• What was done: <a href="https://progress.opensuse.org/issues/73633#note-18">https://progress.opensuse.org/issues/73633#note-18</a></li><li>• How where the changes tested to verify they work: <a href="https://progress.opensuse.org/issues/73633#note-19">https://progress.opensuse.org/issues/73633#note-19</a></li></ul>			
<b>Related issues:</b>			
Related to openQA Infrastructure - action #75016: [osd-admins][alert] Failed ...		<b>Resolved</b>	<b>2020-10-21</b>
Related to openQA Infrastructure - action #75055: grenache-1 can't connect to...		<b>Resolved</b>	<b>2020-10-22</b>
Related to openQA Infrastructure - action #76828: big job queue for ppc as po...		<b>Resolved</b>	<b>2020-10-31</b>
Related to openQA Infrastructure - action #68095: Migrate osd workers from Su...		<b>Resolved</b>	<b>2020-06-15</b>
Related to openQA Infrastructure - action #80128: openqaworker-arm-2 fails to...		<b>Resolved</b>	<b>2020-11-21</b>
Has duplicate openQA Infrastructure - action #77995: worker instances on gren...		<b>Rejected</b>	<b>2020-11-16</b>
Copied to openQA Infrastructure - action #78127: follow-up to #73633 - lesson...		<b>Resolved</b>	

## History

### #1 - 2020-10-21 08:27 - okurz

- Status changed from New to In Progress

- Assignee set to okurz

coolo looked into the issue again this morning, coolo stating "we have a whopping 173 apache slots getting an artefact upload atm, and according to strace they get uploaded in bytes not MBs, [SLES-15-SP3-s390x-63.1@opensuse.org/s390x-kvm-sle15-minimal\\_with\\_sdk63.1\\_installed\\_withhome.qcow2](https://sles-15-sp3-s390x-63.1@opensuse.org/s390x-kvm-sle15-minimal_with_sdk63.1_installed_withhome.qcow2): Processing chunk 520/3037, avg. speed ~19.148 KiB/s [...] the workers have been restarted by salt now - but I stopped the scheduler and so far the 44 jobs running seem to run fine , <https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&from=1603183682771&to=1603199316026> - so the problem started around the time that Nick changed IP routes yesterday, not saying what's cause and what is symptom - but they are surely related [...] So somehow suddenly all workers decided to slowdown uploads ☹☹ [...] So it seems to work all fine again - and all I did was turning it off and on again ☹☹".

since the last problematic incident we have

<https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?orgId=1&refresh=30s&fullscreen&panelId=2&from=now-2d&to=now> and I don't see anything severe showing up there at least. so likely something different? Although I can see that the number of database connections looks *different* at least since 2020-10-20 12:00

The apache response times in

<https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?tab=alert&orgId=1&from=now-2d&to=now&fullscreen&panelId=84&edit> show a significant increase in the response time which we can alert on.

EDIT: Created [https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/384](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/384) for improved monitoring based on apache response time. In the meantime we restarted the openqa-webui service multiple times as well as apache2 and nicksinger removed the manually added IPv6 routes from all machines except grenache-1.

### #2 - 2020-10-21 09:03 - nicksinger

As the problem really escalated yesterday after I enabled a manual IPv6 route and most of OSD's connections where over v6:

```
openqa:~ # ss -tbn4 | wc -l
55
openqa:~ # ss -tbn6 | wc -l
1585
```

I now removed this route from all workers again. The command I used for this was:

```
salt -l error -C 'G@roles:worker' cmd.run 'ip -6 r d default via fe80::1 dev $(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p")'
```

If we see other problems we can think about disabling IPv6 completely for now on the externally connected interfaces like this:

```
salt -l error -C 'G@roles:worker' cmd.run 'echo 1 > /proc/sys/net/ipv6/conf/$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs)/disable_ipv6'
```

### #3 - 2020-10-21 09:07 - nicksinger

We have the initial infra ticket from yesterday about the missing v6 route: <https://infra.nue.suse.com/SelfService/Display.html?id=178626>. In the meantime I stated that all our machines are affected and that we can see severe performance issues over v6. Might be worth to create a new/more explicit one once we're sure we can blame the network.

### #4 - 2020-10-21 09:35 - okurz

From <https://stats.openqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?tab=alert&orgId=1&from=now-3h&to=now&fullscreen&panelId=84&edit> I don't see severe problems right now. I planned to start openqa-scheduler again at 0930Z unless I hear objections.

EDIT:

```
<nicksinger> any objections on disabling v6 on grenache completely? I want to see if it works better then yest
erday with a missing route
<okurz> I suggest we only apply changes one at a time. Do you see severe problems with grenache-1 right now? I
consider it the most important issue that openqa-scheduler is not running so no new jobs will be started
<okurz> started [openqa-scheduler service], btw I hope you guys can all see the annotations in https://stats.o
penqa-monitor.qa.suse.de/d/WebuiDb/webui-summary?tab=alert&orgId=1&from=now-3h&to=now&fullscreen&panelId=84&edit
it ? started, openqa-scheduler on osd again, monitoring the mentioned grafana panel. Updated https://progress.
opensuse.org/issues/73633 and also commented in https://infra.nue.suse.com/SelfService/Display.html?id=178626
. Thanks Nick Singer for the ticket update and the EngInfra ticket reference and making sure that they underst
and the gravity []
```

an alert for "apache response time" is deployed now and it's currently green.

I put the threshold on 500ms avg as I saw that the avg would creep up slowly so I think 1s could give us an alert a bit sooner but still not trigger falsely.

### #5 - 2020-10-21 10:16 - okurz

- Related to action #75016: [osd-admins][alert] Failed systemd services alert (workers): os-autoinst-openvswitch.service (and var-lib-openqa-share.mount) on openqaworker-arm-2 and others added

### #6 - 2020-10-21 13:23 - okurz

- Due date set to 2020-10-23

- Status changed from In Progress to Feedback

For the past hours I was looking into #75016 which I assume to be related. Also I was monitoring grafana alerts (no new alerts during this time) and found no further problems. I am not aware of any current things that do not work. We can try changes regarding "IPv6" again maybe tomorrow as long as no new issues came up or the situation regressed.

### #7 - 2020-10-22 06:38 - nicksinger

- Related to action #75055: grenache-1 can't connect to webui's over IPv4 only added

### #8 - 2020-10-22 20:08 - okurz

- Due date changed from 2020-10-23 to 2020-10-24

osd itself seems to be fine but some machines have problems and are not conducting tests at all. Right now all three arm machines are not conducting tests. On openqaworker-arm-1 which was automatically rebooted (after crash) 5h ago all worker services fail to to reach osd as they try over IPv6 but fail due to the missing route.

What I did now:

```
echo /proc/sys/net/ipv6/conf/$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs)/disab
```

le\_ipv6

```
systemctl restart openqa-worker@\* openqa-worker-cacheservice openqa-worker-cacheservice-minion.service os-aut  
oinst-openvswitch.service
```

and tests start again but this is not persistent.

I guess we could call

```
salt -l error -C 'G@roles:worker' cmd.run 'echo net.ipv6.conf.all.disable_ipv6 = 1 > /etc/sysctl.d/poo73633_de  
bugging.conf && sysctl --load /etc/sysctl.d/poo73633_debugging.conf && systemctl restart openqa-worker@\* open  
qa-worker-cacheservice openqa-worker-cacheservice-minion.service os-autoinst-openvswitch.service'
```

I called that for openqaworker-arm-1 and openqaworker-arm-2 now only. qa-power8-5.qa.suse.de was not reachable and also IPMI SoL gave me nothing so I called power reset and after the machine is up also here like in [#75016](#) the mount point service var-lib-openqa-share.mount failed and I fixed that by restarting with systemctl restart var-lib-openqa-share.mount. I did not remove IPv6 or anything, tests started up but not sure if they will work fine. I can't reach malbec.arch neither ssh nor over IPMI so no progress there.

EDIT: 2020-10-22 21:53: Retrying multiple times I can reach malbec.arch over ipmitool to confirm that "Chassis Power is on" but I can't get it to show anything on SoL so I can only try to trigger a power reset but running something like while [ \$? != 0 ]; do ipmitool -4 -l lanplus -H fsp1-malbec.arch.suse.de -P \$pass power reset && break; done for about 30m on both my computer as well as login1.suse.de fail to establish a session.

EDIT: 2020-10-22 23:40: At a later time I managed to "get through" to malbec and could trigger a power reset. It is conducting tests fine again right now.

EDIT: 2020-10-26 09:24: Applied the same ipv6 disablement from above to grenache-1.qa which failed to run any tests.

### #9 - 2020-10-27 14:07 - nicksinger

So I dug a little more ending up hijacking openqaworker3 as my debugging host. First of, I installed tcpdump to be capable of wireshark tracing over ssh. Nothing too unexpected there besides router advertisements completely missing on the interface for the machine itself. I was however able to spot "Router Solicitation" originating from a QEMU mac (which should only happen if there was a previous RA, so the SUTs can see the router?). I continued probing for all routers (ping ff02::2 - ff02::2 is the multicast address for all routers):

```
64 bytes from fe80::56ab:3aff:fe16:ddc4%eth0: icmp_seq=1 ttl=64 time=0.067 ms  
64 bytes from fe80::56ab:3aff:fe16:dd73%br0: icmp_seq=1 ttl=64 time=0.391 ms (DUP!)  
64 bytes from fe80::56ab:3aff:fe24:358d%br0: icmp_seq=1 ttl=64 time=0.407 ms (DUP!)  
64 bytes from fe80::2e60:cff:fe73:2ac%br0: icmp_seq=1 ttl=64 time=0.422 ms (DUP!)  
64 bytes from fe80::ec4:7aff:fe7a:7896%br0: icmp_seq=1 ttl=64 time=0.471 ms (DUP!)  
64 bytes from fe80::ec4:7aff:fe99:dcd9%br0: icmp_seq=1 ttl=64 time=0.486 ms (DUP!)  
64 bytes from fe80::ec4:7aff:fe43:d6a8%br0: icmp_seq=1 ttl=64 time=0.484 ms (DUP!)  
64 bytes from fe80::fab1:56ff:fed2:7fcf%br0: icmp_seq=1 ttl=64 time=0.500 ms (DUP!)  
64 bytes from fe80::56bf:64ff:fea4:2315%br0: icmp_seq=1 ttl=64 time=0.530 ms (DUP!)  
64 bytes from fe80::6600:6aff:fe73:c434%br0: icmp_seq=1 ttl=64 time=0.529 ms (DUP!)  
64 bytes from fe80::529a:4cff:fe4c:e46d%br0: icmp_seq=1 ttl=64 time=0.554 ms (DUP!)  
64 bytes from fe80::1a03:73ff:fed5:6477%br0: icmp_seq=1 ttl=64 time=0.560 ms (DUP!)  
64 bytes from fe80::9a90:96ff:fea0:fc9b%br0: icmp_seq=1 ttl=64 time=0.569 ms (DUP!)  
64 bytes from fe80::200:5aff:fe9c:4a11%br0: icmp_seq=1 ttl=64 time=0.567 ms (DUP!)  
64 bytes from fe80::3d57:e68f:6817:810f%br0: icmp_seq=1 ttl=64 time=0.579 ms (DUP!)  
64 bytes from fe80::ec4:7aff:fe7a:789e%br0: icmp_seq=1 ttl=64 time=0.587 ms (DUP!)  
64 bytes from fe80::fab1:56ff:febe:b857%br0: icmp_seq=1 ttl=64 time=0.585 ms (DUP!)  
64 bytes from fe80::1a66:daff:fe32:4eec%br0: icmp_seq=1 ttl=64 time=0.602 ms (DUP!)  
64 bytes from fe80::1a66:daff:fe31:9434%br0: icmp_seq=1 ttl=64 time=0.627 ms (DUP!)  
64 bytes from fe80::862b:2bff:feal:28c%br0: icmp_seq=1 ttl=64 time=0.651 ms (DUP!)  
64 bytes from fe80::b002:7eff:fe38:2d23%br0: icmp_seq=1 ttl=64 time=0.660 ms (DUP!)  
64 bytes from fe80::d8a9:36ff:fe86:98b7%br0: icmp_seq=1 ttl=64 time=0.676 ms (DUP!)  
64 bytes from fe80::3617:ebff:fe9e:6902%br0: icmp_seq=1 ttl=64 time=0.757 ms (DUP!)  
64 bytes from fe80::fab1:56ff:feb8:367e%br0: icmp_seq=1 ttl=64 time=1.02 ms (DUP!)  
64 bytes from fe80::2de:fbff:fee3:dafc%br0: icmp_seq=1 ttl=64 time=1.24 ms (DUP!)  
64 bytes from fe80::2de:fbff:fee3:d77c%br0: icmp_seq=1 ttl=64 time=2.84 ms (DUP!)
```

It is very interesting to see so many entries in here. I still need to figure out how exactly how to read this but basically you can see that only one response came from eth0 while all the others came from our bridge on worker3. If all the br0 answers are actually from SUTs is yet unclear to me. But it could show a first problem.

I also found the following which I just leave here for me to parse later:

```
openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'ip -6 neigh'  
openqaworker8.suse.de:  
openqaworker3.suse.de:  
    fe80::1 dev br0 lladdr 00:00:5e:00:02:02 router STALE  
openqaworker9.suse.de:  
    fe80::a3c9:d83f:17aa:8999 dev eth1 lladdr d4:81:d7:5a:a3:9c STALE  
    fe80::36ac:19a7:3193:7081 dev eth1 lladdr 0a:00:00:00:00:33 STALE  
    fe80::216:3eff:fe48:17ff dev eth1 lladdr 00:16:3e:48:17:ff STALE
```

```
fe80::5054:ff:fe44:d766 dev eth1 lladdr 52:54:00:44:d7:66 STALE
fe80::5054:ff:fe44:d765 dev eth1 lladdr 52:54:00:44:d7:65 STALE
fe80::5054:ff:fe44:d768 dev eth1 lladdr 52:54:00:44:d7:68 STALE
fe80::5054:ff:fe44:d767 dev eth1 lladdr 52:54:00:44:d7:67 STALE
fe80::6600:6aff:fe75:72 dev eth1 lladdr 64:00:6a:75:00:72 STALE
fe80::c3ab:62d0:2723:6249 dev eth1 lladdr 64:00:6a:75:00:72 STALE
2620:113:80c0:8080::4 dev eth1 FAILED
fe80::501:abb4:eb5c:6686 dev eth1 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::5054:ff:fe30:a4d9 dev eth1 lladdr 52:54:00:30:a4:d9 STALE
fe80::208:2fff:feed:8f15 dev eth1 lladdr 00:08:02:ed:8f:15 STALE
fe80::2af1:eff:fe41:cef3 dev eth1 lladdr 28:f1:0e:41:ce:f3 STALE
fe80::1 dev eth1 lladdr 00:00:5e:00:02:02 router STALE
fe80::ec4:7aff:fe7a:7736 dev eth1 lladdr 0c:c4:7a:7a:77:36 STALE
fe80::4950:d671:f08c:c9c3 dev eth1 lladdr 18:db:f2:46:1e:1d STALE
fe80::9249:faff:fe06:82d8 dev eth1 lladdr 90:49:fa:06:82:d8 STALE
fe80::2de:fbff:fee3:d77c dev eth1 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::d681:d7ff:fe5a:a39c dev eth1 lladdr d4:81:d7:5a:a3:9c STALE
fe80::800:ff:fe00:15 dev eth1 lladdr 0a:00:00:00:00:15 STALE
fe80::2de:fbff:fee3:dafc dev eth1 lladdr 00:de:fb:e3:da:fc router STALE
fe80::56ab:3aff:fe16:ddc4 dev eth1 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::5054:ff:fe29:137f dev eth1 lladdr 52:54:00:29:13:7f STALE
fe80::1a66:daff:fe00:bbaa dev eth1 lladdr 18:66:da:00:bb:aa STALE
fe80::800:ff:fe00:32 dev eth1 lladdr 0a:00:00:00:00:32 STALE
fe80::5054:ff:fef4:ecb8 dev eth1 lladdr 52:54:00:f4:ec:b8 STALE
fe80::5054:ff:fe87:8cc4 dev eth1 lladdr 52:54:00:87:8c:c4 STALE
openqaworker6.suse.de:
fe80::5054:ff:fe44:d767 dev eth0 lladdr 52:54:00:44:d7:67 STALE
fe80::1 dev eth0 lladdr 00:00:5e:00:02:02 router STALE
fe80::2de:fbff:fee3:dafc dev eth0 lladdr 00:de:fb:e3:da:fc router STALE
fe80::1a66:daff:fe00:bbaa dev eth0 lladdr 18:66:da:00:bb:aa STALE
fe80::800:ff:fe00:15 dev eth0 lladdr 0a:00:00:00:00:15 STALE
fe80::9249:faff:fe06:82d8 dev eth0 lladdr 90:49:fa:06:82:d8 STALE
fe80::56ab:3aff:fe16:ddc4 dev eth0 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::5054:ff:fe44:d765 dev eth0 lladdr 52:54:00:44:d7:65 STALE
fe80::d681:d7ff:fe5a:a39c dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::216:3eff:fe48:17ff dev eth0 lladdr 00:16:3e:48:17:ff STALE
fe80::208:2fff:feed:8f15 dev eth0 lladdr 00:08:02:ed:8f:15 STALE
fe80::800:ff:fe00:32 dev eth0 lladdr 0a:00:00:00:00:32 STALE
fe80::6600:6aff:fe75:72 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::5054:ff:fe30:a4d9 dev eth0 lladdr 52:54:00:30:a4:d9 STALE
fe80::5054:ff:fe44:d768 dev eth0 lladdr 52:54:00:44:d7:68 STALE
fe80::36ac:19a7:3193:7081 dev eth0 lladdr 0a:00:00:00:00:33 STALE
fe80::ec4:7aff:fe7a:7736 dev eth0 lladdr 0c:c4:7a:7a:77:36 STALE
fe80::2908:884f:5368:dda dev eth0 lladdr c8:f7:50:40:f4:69 STALE
fe80::2af1:eff:fe41:cef3 dev eth0 lladdr 28:f1:0e:41:ce:f3 STALE
fe80::5054:ff:fe87:8cc4 dev eth0 lladdr 52:54:00:87:8c:c4 STALE
fe80::5054:ff:fe44:d766 dev eth0 lladdr 52:54:00:44:d7:66 STALE
fe80::5054:ff:feb1:4de dev eth0 lladdr 52:54:00:b1:04:de STALE
fe80::501:abb4:eb5c:6686 dev eth0 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::2de:fbff:fee3:d77c dev eth0 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::a3c9:d83f:17aa:8999 dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::5054:ff:fef4:ecb8 dev eth0 lladdr 52:54:00:f4:ec:b8 STALE
fe80::4950:d671:f08c:c9c3 dev eth0 lladdr 18:db:f2:46:1e:1d STALE
fe80::c3ab:62d0:2723:6249 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::5054:ff:fe29:137f dev eth0 lladdr 52:54:00:29:13:7f STALE
QA-Power8-4-kvm.qa.suse.de:
fe80::1 dev eth3 lladdr 00:00:5e:00:02:04 router STALE
fe80::f46b:41ff:feb7:9502 dev eth3 lladdr f6:6b:41:b7:95:02 STALE
fe80::2de:fbff:fee3:dafc dev eth3 lladdr 00:de:fb:e3:da:fc router STALE
fe80::215:5dff:fe43:a241 dev eth3 lladdr 00:15:5d:43:a2:41 STALE
fe80::2de:fbff:fee3:d77c dev eth3 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::f46b:44ff:fe50:f502 dev eth3 lladdr f6:6b:44:50:f5:02 STALE
fe80::5054:ff:fe47:10e4 dev eth3 lladdr 52:54:00:47:10:e4 STALE
fe80::216:3eff:fe32:3671 dev eth3 lladdr 00:16:3e:32:36:71 STALE
fe80::216:3eff:fec3:d305 dev eth3 lladdr 00:16:3e:c3:d3:05 STALE
fe80::f46b:45ff:fe75:7e02 dev eth3 lladdr f6:6b:45:75:7e:02 STALE
fe80::ae1f:6bff:fe01:130 dev eth3 lladdr ac:1f:6b:01:01:30 STALE
fe80::f46b:47ff:fe57:de02 dev eth3 lladdr f6:6b:47:57:de:02 STALE
fe80::215:5dff:fe43:a23d dev eth3 lladdr 00:15:5d:43:a2:3d STALE
fe80::dc86:c1ff:fe33:d97f dev eth3 lladdr de:86:c1:33:d9:7f STALE
fe80::1e1b:dff:feef:735c dev eth3 lladdr 1c:1b:0d:ef:73:5c STALE
fe80::216:3eff:fe32:6543 dev eth3 lladdr 00:16:3e:32:65:43 STALE
fe80::e2d5:5eff:fea7:e824 dev eth3 lladdr e0:d5:5e:a7:e8:24 STALE
fe80::215:5dff:fe43:a23b dev eth3 lladdr 00:15:5d:43:a2:3b STALE
fe80::f46b:4aff:fef5:d602 dev eth3 lladdr f6:6b:4a:f5:d6:02 STALE
```

```
fe80::215:5dff:fe43:a239 dev eth3 lladdr 00:15:5d:43:a2:39 STALE
fe80::f46b:46ff:fe0a:3202 dev eth3 lladdr f6:6b:46:0a:32:02 STALE
fe80::216:3eff:fe32:8923 dev eth3 lladdr 00:16:3e:32:89:23 STALE
fe80::f46b:4fff:fe78:3902 dev eth3 lladdr f6:6b:4f:78:39:02 STALE
fe80::5054:ff:fea2:abb2 dev eth3 lladdr 52:54:00:a2:ab:b2 STALE
fe80::20c:29ff:fe20:339f dev eth3 lladdr 00:0c:29:20:33:9f STALE
fe80::225:90ff:fe9a:cb5e dev eth3 lladdr 00:25:90:9a:cb:5e STALE
fe80::423:f5ff:fe3c:2c73 dev eth3 lladdr 06:23:f5:3c:2c:73 STALE
fe80::f46b:43ff:fed5:9d02 dev eth3 lladdr f6:6b:43:d5:9d:02 STALE
fe80::ff:feel:a5b4 dev eth3 lladdr 02:00:00:e1:a5:b4 STALE
fe80::5054:ff:fe40:4ale dev eth3 lladdr 52:54:00:40:4a:1e STALE
fe80::ec4:7aff:fe6c:400a dev eth3 lladdr 0c:c4:7a:6c:40:0a STALE
fe80::215:5dff:fe43:a23e dev eth3 lladdr 00:15:5d:43:a2:3e STALE
fe80::f46b:45ff:fee9:d803 dev eth3 lladdr f6:6b:45:e9:d8:03 STALE
fe80::215:5dff:fe43:a23c dev eth3 lladdr 00:15:5d:43:a2:3c STALE
fe80::ff:fee0:a4b3 dev eth3 lladdr 02:00:00:e0:a4:b3 STALE
fe80::5054:ff:fe55:613f dev eth3 lladdr 52:54:00:55:61:3f STALE
fe80::20c:29ff:fe9d:6297 dev eth3 lladdr 00:0c:29:9d:62:97 STALE
openqaworker2.suse.de:
fe80::5054:ff:fe30:a4d9 dev br0 lladdr 52:54:00:30:a4:d9 STALE
fe80::4950:d671:f08c:c9c3 dev br0 lladdr 18:db:f2:46:1e:1d STALE
fe80::2de:fbff:fee3:dafc dev br0 lladdr 00:de:fb:e3:da:fc router STALE
fe80::ec4:7aff:fe7a:7736 dev br0 lladdr 0c:c4:7a:7a:77:36 STALE
fe80::6600:6aff:fe75:72 dev br0 lladdr 64:00:6a:75:00:72 STALE
fe80::5054:ff:fe29:137f dev br0 lladdr 52:54:00:29:13:7f STALE
fe80::800:ff:fe00:15 dev br0 lladdr 0a:00:00:00:00:15 STALE
fe80::56ab:3aff:fe16:ddc4 dev br0 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::9249:faff:fe06:82d8 dev br0 lladdr 90:49:fa:06:82:d8 STALE
fe80::1 dev br0 lladdr 00:00:5e:00:02:02 router STALE
fe80::2af1:eff:fe41:cef3 dev br0 lladdr 28:f1:0e:41:ce:f3 STALE
2620:113:80c0:8080::5 dev br0 FAILED
fe80::a3c9:d83f:17aa:8999 dev br0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::d681:d7ff:fe5a:a39c dev br0 lladdr d4:81:d7:5a:a3:9c STALE
2620:113:80c0:8080::4 dev br0 FAILED
fe80::5054:ff:fef4:ecb8 dev br0 lladdr 52:54:00:f4:ec:b8 STALE
fe80::208:2ff:feed:8f15 dev br0 lladdr 00:08:02:ed:8f:15 STALE
fe80::2de:fbff:fee3:d77c dev br0 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::5054:ff:fe44:d768 dev br0 lladdr 52:54:00:44:d7:68 STALE
fe80::501:abb4:eb5c:6686 dev br0 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::5054:ff:fe44:d767 dev br0 lladdr 52:54:00:44:d7:67 STALE
fe80::5054:ff:fe87:8cc4 dev br0 lladdr 52:54:00:87:8c:c4 STALE
fe80::c3ab:62d0:2723:6249 dev br0 lladdr 64:00:6a:75:00:72 STALE
fe80::5054:ff:fe44:d766 dev br0 lladdr 52:54:00:44:d7:66 STALE
fe80::800:ff:fe00:32 dev br0 lladdr 0a:00:00:00:00:32 STALE
fe80::1a66:daff:fe00:bbaa dev br0 lladdr 18:66:da:00:bb:aa STALE
fe80::36ac:19a7:3193:7081 dev br0 lladdr 0a:00:00:00:00:33 STALE
fe80::216:3eff:fe48:17ff dev br0 lladdr 00:16:3e:48:17:ff STALE
fe80::5054:ff:fe44:d765 dev br0 lladdr 52:54:00:44:d7:65 STALE
openqaworker5.suse.de:
fe80::5054:ff:fe30:a4d9 dev eth0 lladdr 52:54:00:30:a4:d9 STALE
fe80::6600:6aff:fe75:72 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::208:2ff:feed:8f15 dev eth0 lladdr 00:08:02:ed:8f:15 STALE
fe80::4950:d671:f08c:c9c3 dev eth0 lladdr 18:db:f2:46:1e:1d STALE
fe80::5054:ff:fe44:d765 dev eth0 lladdr 52:54:00:44:d7:65 STALE
fe80::800:ff:fe00:32 dev eth0 lladdr 0a:00:00:00:00:32 STALE
fe80::9249:faff:fe06:82d8 dev eth0 lladdr 90:49:fa:06:82:d8 STALE
fe80::d681:d7ff:fe5a:a39c dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::216:3eff:fe48:17ff dev eth0 lladdr 00:16:3e:48:17:ff STALE
fe80::1a66:daff:fe00:bbaa dev eth0 lladdr 18:66:da:00:bb:aa STALE
fe80::5054:ff:fe44:d767 dev eth0 lladdr 52:54:00:44:d7:67 STALE
fe80::a3c9:d83f:17aa:8999 dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::5054:ff:fe29:137f dev eth0 lladdr 52:54:00:29:13:7f STALE
fe80::36ac:19a7:3193:7081 dev eth0 lladdr 0a:00:00:00:00:33 STALE
fe80::2af1:eff:fe41:cef3 dev eth0 lladdr 28:f1:0e:41:ce:f3 STALE
fe80::5054:ff:fe87:8cc4 dev eth0 lladdr 52:54:00:87:8c:c4 STALE
fe80::56ab:3aff:fe16:ddc4 dev eth0 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::501:abb4:eb5c:6686 dev eth0 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::5054:ff:fe44:d766 dev eth0 lladdr 52:54:00:44:d7:66 STALE
fe80::5054:ff:feb1:4de dev eth0 lladdr 52:54:00:b1:04:de STALE
fe80::5054:ff:fef4:ecb8 dev eth0 lladdr 52:54:00:f4:ec:b8 STALE
fe80::1 dev eth0 lladdr 00:00:5e:00:02:02 router STALE
2620:113:80c0:8080::4 dev eth0 FAILED
fe80::800:ff:fe00:15 dev eth0 lladdr 0a:00:00:00:00:15 STALE
fe80::2de:fbff:fee3:dafc dev eth0 lladdr 00:de:fb:e3:da:fc router STALE
fe80::c3ab:62d0:2723:6249 dev eth0 lladdr 64:00:6a:75:00:72 STALE
```

```

fe80::5054:ff:fe44:d768 dev eth0 lladdr 52:54:00:44:d7:68 STALE
fe80::2de:fbff:fee3:d77c dev eth0 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::ec4:7aff:fe7a:7736 dev eth0 lladdr 0c:c4:7a:7a:77:36 STALE
fe80::2908:884f:5368:dda dev eth0 lladdr c8:f7:50:40:f4:69 STALE
grenache-1.qa.suse.de:
openqaworker10.suse.de:
fe80::c3ab:62d0:2723:6249 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::800:ff:fe00:15 dev eth0 lladdr 0a:00:00:00:00:15 STALE
fe80::5054:ff:fe44:d768 dev eth0 lladdr 52:54:00:44:d7:68 STALE
fe80::4950:d671:f08c:c9c3 dev eth0 lladdr 18:db:f2:46:1e:1d STALE
fe80::501:abb4:eb5c:6686 dev eth0 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::800:ff:fe00:32 dev eth0 lladdr 0a:00:00:00:00:32 STALE
fe80::2de:fbff:fee3:dafc dev eth0 lladdr 00:de:fb:e3:da:fc router STALE
fe80::208:2ff:feed:8f15 dev eth0 lladdr 00:08:02:ed:8f:15 STALE
2620:113:80c0:8080::5 dev eth0 FAILED
fe80::9249:faff:fe06:82d8 dev eth0 lladdr 90:49:fa:06:82:d8 STALE
fe80::1 dev eth0 lladdr 00:00:5e:00:02:02 router STALE
fe80::5054:ff:fe44:d765 dev eth0 lladdr 52:54:00:44:d7:65 STALE
fe80::1a66:daff:fe00:bbaa dev eth0 lladdr 18:66:da:00:bb:aa STALE
fe80::6600:6aff:fe75:72 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::5054:ff:fe44:d767 dev eth0 lladdr 52:54:00:44:d7:67 STALE
fe80::5054:ff:fe44:ecb8 dev eth0 lladdr 52:54:00:f4:ec:b8 STALE
fe80::5054:ff:fe29:137f dev eth0 lladdr 52:54:00:29:13:7f STALE
fe80::20d:b9ff:fe01:ea8 dev gre_sys FAILED
fe80::2de:fbff:fee3:d77c dev eth0 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::a3c9:d83f:17aa:8999 dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::5054:ff:fe30:a4d9 dev eth0 lladdr 52:54:00:30:a4:d9 STALE
fe80::d681:d7ff:fe5a:a39c dev eth0 lladdr d4:81:d7:5a:a3:9c STALE
fe80::216:3eff:fe48:17ff dev eth0 lladdr 00:16:3e:48:17:ff STALE
fe80::5054:ff:fe87:8cc4 dev eth0 lladdr 52:54:00:87:8c:c4 STALE
fe80::56ab:3aff:fe16:ddc4 dev eth0 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::2af1:eff:fe41:cef3 dev eth0 lladdr 28:f1:0e:41:ce:f3 STALE
fe80::36ac:19a7:3193:7081 dev eth0 lladdr 0a:00:00:00:00:33 STALE
fe80::5054:ff:fe44:d766 dev eth0 lladdr 52:54:00:44:d7:66 STALE
fe80::ec4:7aff:fe7a:7736 dev eth0 lladdr 0c:c4:7a:7a:77:36 STALE
openqaworker13.suse.de:
fe80::5054:ff:fe44:d766 dev eth0 lladdr 52:54:00:44:d7:66 STALE
fe80::1 dev eth0 lladdr 00:00:5e:00:02:02 router STALE
fe80::800:ff:fe00:32 dev eth0 lladdr 0a:00:00:00:00:32 STALE
fe80::6600:6aff:fe75:72 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::2af1:eff:fe41:cef3 dev eth0 lladdr 28:f1:0e:41:ce:f3 STALE
fe80::c3ab:62d0:2723:6249 dev eth0 lladdr 64:00:6a:75:00:72 STALE
fe80::36ac:19a7:3193:7081 dev eth0 lladdr 0a:00:00:00:00:33 STALE
fe80::ec4:7aff:fe7a:7736 dev eth0 lladdr 0c:c4:7a:7a:77:36 STALE
fe80::5054:ff:fe44:d767 dev eth0 lladdr 52:54:00:44:d7:67 STALE
fe80::216:3eff:fe48:17ff dev eth0 lladdr 00:16:3e:48:17:ff STALE
2620:113:80c0:8080::5 dev eth0 FAILED
fe80::9249:faff:fe06:82d8 dev eth0 lladdr 90:49:fa:06:82:d8 STALE
fe80::5054:ff:fe44:d768 dev eth0 lladdr 52:54:00:44:d7:68 STALE
fe80::5054:ff:fe30:a4d9 dev eth0 lladdr 52:54:00:30:a4:d9 STALE
fe80::5054:ff:fe44:ecb8 dev eth0 lladdr 52:54:00:f4:ec:b8 STALE
fe80::5054:ff:fe29:137f dev eth0 lladdr 52:54:00:29:13:7f STALE
fe80::2de:fbff:fee3:dafc dev eth0 lladdr 00:de:fb:e3:da:fc router STALE
fe80::4950:d671:f08c:c9c3 dev eth0 lladdr 18:db:f2:46:1e:1d STALE
fe80::5054:ff:fe44:d765 dev eth0 lladdr 52:54:00:44:d7:65 STALE
fe80::208:2ff:feed:8f15 dev eth0 lladdr 00:08:02:ed:8f:15 STALE
fe80::800:ff:fe00:15 dev eth0 lladdr 0a:00:00:00:00:15 STALE
fe80::56ab:3aff:fe16:ddc4 dev eth0 lladdr 54:ab:3a:16:dd:c4 router STALE
fe80::501:abb4:eb5c:6686 dev eth0 lladdr e4:b9:7a:e4:aa:ad STALE
fe80::5054:ff:fe87:8cc4 dev eth0 lladdr 52:54:00:87:8c:c4 STALE
fe80::2de:fbff:fee3:d77c dev eth0 lladdr 00:de:fb:e3:d7:7c router STALE
fe80::1a66:daff:fe00:bbaa dev eth0 lladdr 18:66:da:00:bb:aa STALE
openqaworker-arm-1.suse.de:
openqaworker-arm-2.suse.de:
QA-Power8-5-kvm.qa.suse.de:
Minion did not return. [Not connected]
malbec.arch.suse.de:
Minion did not return. [Not connected]
openqaworker-arm-3.suse.de:
Minion did not return. [Not connected]

```

#### #10 - 2020-10-30 05:57 - okurz

nicksinger in <https://infra.nue.suse.com/SelfService/Display.html?id=178626> mmaher asked the question "Did the operation with the s390 host in the qa network helped in this issue? is it still the case? or any other news?". *Something* is certainly still wrong but I think what we could do is to provide

"steps to reproduce" in EngInfra tickets. Otherwise the poor lads and lassies really do not have a better chance then to ask the reporter "is it still happening". And here I am not even super sure. So is the way to test: "Reboot worker machine, make sure no workaround disables IPv6 and call ping6 -c 1 www.opensuse.org to check if IPv6 works?" or is ping6 -c 1 openqa.suse.de enough?

#### #11 - 2020-10-31 09:09 - okurz

- Related to action #76828: big job queue for ppc as powerqaworker-qam-1.qa and malbec.arch and qa-power8-5-kvm were not active added

#### #12 - 2020-11-02 12:20 - nicksinger

okurz wrote:

[nicksinger](https://infra.nue.suse.com/SelfService/Display.html?id=178626) in <https://infra.nue.suse.com/SelfService/Display.html?id=178626> mmaher asked the question "Did the operation with the s390 host in the qa network helped in this issue? is it still the case? or any other news?". *Something* is certainly still wrong but I think what we could do is to provide "steps to reproduce" in EngInfra tickets. Otherwise the poor lads and lassies really do not have a better chance then to ask the reporter "is it still happening". And here I am not even super sure. So is the way to test: "Reboot worker machine, make sure no workaround disables IPv6 and call ping6 -c 1 www.opensuse.org to check if IPv6 works?" or is ping6 -c 1 openqa.suse.de enough?

strictly speaking about the v6 issue I think your first approach is the best. It should also be possible to do all this "at runtime" but safest is a reboot of course.

#### #13 - 2020-11-04 14:07 - nicksinger

the repair of powerqaworker-qam-1 showed some interesting results as the machine was broken long enough to not get the most recent salt updates. Right after the machine was started:

```
powerqaworker-qam-1:~ # ip -6 r s
2620:113:80c0:80a0::/64 dev eth4 proto kernel metric 256 expires 3535sec pref medium
fe80::/64 dev br1 proto kernel metric 256 pref medium
fe80::/64 dev eth4 proto kernel metric 256 pref medium
default via fe80::1 dev eth4 proto ra metric 1024 expires 1735sec hoplimit 64 pref medium
```

at this time, the salt-key was blocklisted and therefore no states where applied. To conclude my work on <https://progress.opensuse.org/issues/68053> I accepted the salt-key on OSD once again and issues an manual "state.highstate". Here is what was changed:

```
openqa:~ # salt 'powerqaworker-qam-1' state.highstate
powerqaworker-qam-1:
-----
      ID: firewallld
      Function: service.running
      Result: True
      Comment: Service firewallld is already enabled, and is running
      Started: 14:57:32.501786
      Duration: 545.143 ms
      Changes:
  -----
          firewallld:
              True
  -----
      ID: grub-conf
      Function: augeas.change
      Result: True
      Comment: Changes have been saved
      Started: 14:57:37.141460
      Duration: 176.998 ms
      Changes:
  -----
          diff:
              ---
              +++
              @@ -14 +14 @@
              -GRUB_CMDLINE_LINUX_DEFAULT="nospec kvm.nested=1 kvm_intel.nested=1 kvm_amd.nested=1 kvm-arm
.nested=1 crashkernel=210M"
              +GRUB_CMDLINE_LINUX_DEFAULT=" nospec kvm.nested=1 kvm_intel.nested=1 kvm_amd.nested=1 kvm-ar
m.nested=1 crashkernel=210M"
  -----
      ID: grub2-mkconfig > /boot/grub2/grub.cfg
      Function: cmd.run
      Result: True
      Comment: Command "grub2-mkconfig > /boot/grub2/grub.cfg" run
      Started: 14:57:37.321017
      Duration: 708.689 ms
      Changes:
  -----
```

```
pid:
  30665
retcode:
  0
stderr:
  Generating grub configuration file ...
  Found linux image: /boot/vmlinuz-4.12.14-lp151.28.75-default
  Found initrd image: /boot/initrd-4.12.14-lp151.28.75-default
  Found linux image: /boot/vmlinuz-4.12.14-lp151.28.48-default
  Found initrd image: /boot/initrd-4.12.14-lp151.28.48-default
  done
stdout:
```

```
-----
ID: telegraf
Function: service.running
Result: True
Comment: Started Service telegraf
Started: 14:57:38.276106
Duration: 171.584 ms
Changes:
-----
telegraf:
  True
```

Summary for powerqaworker-qam-1

```
-----
Succeeded: 270 (changed=4)
Failed:    0
-----
Total states run:    270
Total run time:    35.355 s
```

and afterwards:

```
powerqaworker-qam-1:~ # ip -6 r s
2620:113:80c0:80a0::/64 dev eth4 proto kernel metric 256 expires 3355sec pref medium
fe80::/64 dev br1 proto kernel metric 256 pref medium
fe80::/64 dev eth4 proto kernel metric 256 pref medium
```

so everything points to firewalld ATM. Disabling firewalld didn't bring the default route back. Will see if I can somehow restore a "working system" again to bisect where our firewalld behaves wrong.

#### #14 - 2020-11-04 15:01 - nicksinger

- File `ip6tables-save.firewalld.txt` added
- File `ip6tables-save.susefirewall.txt` added

firewalld is certainly to blame here. I've collected the dump of ip6tables but that's too much for me to digest for today  
EDIT: colored diff of these two files can be found at <https://w3.suse.de/~nsinger/diff.html>

#### #15 - 2020-11-05 05:46 - okurz

- Status changed from *Feedback* to *In Progress*
- Assignee changed from *okurz* to *nicksinger*

Great news. Please continue the firewalld investigation.

#### #16 - 2020-11-06 11:13 - nicksinger

seems like firewalld was just the trigger. Currently following the hint that if `net.ipv6.conf.all.forwarding = 1` is set then `net.ipv6.conf.eth1.accept_ra` needs to be set to 2 to accept RA's which seem to be the base for `wicked-dhcp6`

#### #17 - 2020-11-06 12:03 - nicksinger

alright so my suspicion was confirmed. Something caused `net.ipv6.conf.all.forwarding` to be set to 1 - I *assume* this was implicitly done by firewalld. According to <https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt> :

```
accept_ra - INTEGER
  Accept Router Advertisements; autoconfigure using them.
```

```
It also determines whether or not to transmit Router
Solicitations. If and only if the functional setting is to
accept Router Advertisements, Router Solicitations will be
```



transmitted.

Possible values are:

- 0 Do not accept Router Advertisements.
- 1 Accept Router Advertisements if forwarding is disabled.
- 2 Overrule forwarding behaviour. Accept Router Advertisements even if forwarding is enabled.

Functional default: enabled if local forwarding is disabled.  
disabled if local forwarding is enabled.

Therefore our workers didn't receive any RA from the NEXUS anymore resulting in dhcpv6 (from wicked) not being able to configure IPv6 properly any longer. That's why we saw proper configured link-local addresses (fe80::/64) but no link-global (2620:113:80c0:8080::/64 - this is the suse prefix). Also the default route over fe80::1 was missing because of the missing (or rather, not accepted) RA's.

BTW: I was able to reproduce the severe performance impact that we saw once we added fe80::1 manually as default route. This is caused if you *only* have a default route but no route for your prefix resulting in ICMP redirects from the router each and every time the machine tries to reach something in its own v6-subnet (which is basically *every* machine inside SUSE). This redirect resulted in a *massive* amount of re-transmitted TCP packages dropping the performance down to max 5MB/s and even stall connections for almost the rest of the time.

#### #18 - 2020-11-06 12:15 - nicksinger

This was the current (broken) state. Please note that worker8, qam-1, worker2 and both arms where my test subjects so it's expected to look correct there. All others show no default route for v6:

```
openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'ip -6 r s | grep default'
openqaworker3.suse.de:
openqaworker9.suse.de:
openqaworker8.suse.de:
  default via fe80::1 dev eth1 proto ra metric 1024 expires 3418sec hoplimit 64 pref medium
openqaworker6.suse.de:
QA-Power8-4-kvm.qa.suse.de:
powerqaworker-qam-1:
  default via fe80::1 dev eth4 metric 1024 pref medium
openqaworker5.suse.de:
QA-Power8-5-kvm.qa.suse.de:
openqaworker2.suse.de:
  default via fe80::1 dev br0 proto ra metric 1024 expires 3418sec hoplimit 64 pref medium
malbec.arch.suse.de:
grenache-1.qa.suse.de:
openqaworker13.suse.de:
openqaworker10.suse.de:
openqaworker-arm-1.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3417sec hoplimit 64 pref medium
openqaworker-arm-2.suse.de:
  default via fe80::1 dev eth1 proto ra metric 1024 expires 3417sec hoplimit 64 pref medium
```

So what I did now to fix this is the following:

1. net.ipv6.conf.all.disable\_ipv6=0 to enable ipv6 on all interfaces again removing any previous workaround on the machines
2. With `$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs)` I get the default interface for v4 traffic. Since we use the same interface for both address types we can just use it as default for all v6 operations that follow now
3. `sysctl net.ipv6.conf.$default_interface.disable_ipv6=1` disable v6 explicitly on the uplink so we can see if it worked afterwards
4. `sysctl net.ipv6.conf.$default_interface.accept_ra=2` enable RA's on the uplink only. We could set it for *all* interfaces but SUTs could misbehave and shouldn't affect the workers interface...
5. `sysctl net.ipv6.conf.$default_interface.disable_ipv6=0` bring back v6 to instantly trigger a SLAAC

The actual salt command looks a little messy but basically the steps described above:

```
openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'sysctl net.ipv6.conf.all.disable_ipv6=0; sysctl net.ipv6
.conf.$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs).disable_ipv6=1; sysctl net.i
pv6.conf.$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs).accept_ra=2; sysctl net.i
pv6.conf.$(ip r s | grep default | sed -n "s/^.*dev \(.*\) proto dhcp/\1/p" | xargs).disable_ipv6=0'
openqaworker8.suse.de:
  net.ipv6.conf.all.disable_ipv6 = 0
  net.ipv6.conf.eth1.disable_ipv6 = 1
  net.ipv6.conf.eth1.accept_ra = 2
  net.ipv6.conf.eth1.disable_ipv6 = 0
openqaworker3.suse.de:
  net.ipv6.conf.all.disable_ipv6 = 0
  net.ipv6.conf.br0.disable_ipv6 = 1
  net.ipv6.conf.br0.accept_ra = 2
  net.ipv6.conf.br0.disable_ipv6 = 0
powerqaworker-qam-1:
  net.ipv6.conf.all.disable_ipv6 = 0
```

```

net.ipv6.conf.eth4.disable_ipv6 = 1
net.ipv6.conf.eth4.accept_ra = 2
net.ipv6.conf.eth4.disable_ipv6 = 0
QA-Power8-5-kvm.qa.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth3.disable_ipv6 = 1
net.ipv6.conf.eth3.accept_ra = 2
net.ipv6.conf.eth3.disable_ipv6 = 0
QA-Power8-4-kvm.qa.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth3.disable_ipv6 = 1
net.ipv6.conf.eth3.accept_ra = 2
net.ipv6.conf.eth3.disable_ipv6 = 0
malbec.arch.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth4.disable_ipv6 = 1
net.ipv6.conf.eth4.accept_ra = 2
net.ipv6.conf.eth4.disable_ipv6 = 0
grenache-1.qa.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0
openqaworker6.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0
openqaworker9.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth1.disable_ipv6 = 1
net.ipv6.conf.eth1.accept_ra = 2
net.ipv6.conf.eth1.disable_ipv6 = 0
openqaworker-arm-1.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0
openqaworker13.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0
openqaworker5.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0
openqaworker-arm-2.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth1.disable_ipv6 = 1
net.ipv6.conf.eth1.accept_ra = 2
net.ipv6.conf.eth1.disable_ipv6 = 0
openqaworker2.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.br0.disable_ipv6 = 1
net.ipv6.conf.br0.accept_ra = 2
net.ipv6.conf.br0.disable_ipv6 = 0
openqaworker10.suse.de:
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.eth0.disable_ipv6 = 1
net.ipv6.conf.eth0.accept_ra = 2
net.ipv6.conf.eth0.disable_ipv6 = 0

```

After I issued the command from above:

```

openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'ip -6 r s | grep default'
openqaworker3.suse.de:
  default via fe80::1 dev br0 proto ra metric 1024 expires 3491sec hoplimit 64 pref medium
openqaworker8.suse.de:
  default via fe80::1 dev eth1 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium
openqaworker5.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium
openqaworker9.suse.de:
  default via fe80::1 dev eth1 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium

```

```

openqaworker2.suse.de:
  default via fe80::1 dev br0 proto ra metric 1024 expires 3494sec hoplimit 64 pref medium
QA-Power8-5-kvm.qa.suse.de:
  default via fe80::1 dev eth3 proto ra metric 1024 expires 1691sec hoplimit 64 pref medium
openqaworker6.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium
powerqaworker-qam-1:
  default via fe80::1 dev eth4 proto ra metric 1024 expires 1692sec hoplimit 64 pref medium
QA-Power8-4-kvm.qa.suse.de:
  default via fe80::1 dev eth3 proto ra metric 1024 expires 1690sec hoplimit 64 pref medium
malbec.arch.suse.de:
  default via fe80::1 dev eth4 proto ra metric 1024 expires 3502sec hoplimit 64 pref medium
grenache-1.qa.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 1691sec hoplimit 64 pref medium
openqaworker10.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium
openqaworker13.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium
openqaworker-arm-1.suse.de:
  default via fe80::1 dev eth0 proto ra metric 1024 expires 3492sec hoplimit 64 pref medium
openqaworker-arm-2.suse.de:
  default via fe80::1 dev eth1 proto ra metric 1024 expires 3493sec hoplimit 64 pref medium

```

### #19 - 2020-11-06 12:53 - nicksinger

After applying these changes, OSD can be reached over v6 from all machines:

```

openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'ping6 -c 1 openqa.suse.de'
openqaworker2.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.281 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.281/0.281/0.281/0.000 ms
openqaworker8.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.664 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.664/0.664/0.664/0.000 ms
openqaworker3.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.496 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.496/0.496/0.496/0.000 ms
openqaworker6.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.167 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.167/0.167/0.167/0.000 ms
openqaworker9.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.381 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.381/0.381/0.381/0.000 ms
QA-Power8-5-kvm.qa.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=63 time=0.278 ms

  --- openqa.suse.de ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.278/0.278/0.278/0.000 ms
openqaworker5.suse.de:
  PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
  64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.614 ms

  --- openqa.suse.de ping statistics ---

```

```

1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.614/0.614/0.614/0.000 ms
powerqaworker-qam-1:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=63 time=0.214 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.214/0.214/0.214/0.000 ms
QA-Power8-4-kvm.qa.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=63 time=0.197 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.197/0.197/0.197/0.000 ms
malbec.arch.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=63 time=0.183 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.183/0.183/0.183/0.000 ms
grenache-1.qa.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=63 time=0.478 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.478/0.478/0.478/0.000 ms
openqaworker10.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.154 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.154/0.154/0.154/0.000 ms
openqaworker13.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.236 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.236/0.236/0.236/0.000 ms
openqaworker-arm-1.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=0.297 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.297/0.297/0.297/0.000 ms
openqaworker-arm-2.suse.de:
PING openqa.suse.de(openqa.suse.de (2620:113:80c0:8080:10:160:0:207)) 56 data bytes
64 bytes from openqa.suse.de (2620:113:80c0:8080:10:160:0:207): icmp_seq=1 ttl=64 time=3.09 ms

--- openqa.suse.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 3.090/3.090/3.090/0.000 ms

```

Just because we saw performance issues with the last workaround I deployed I wanted to have a speedtest too. The other side was running on my workstation which is in the same VLAN but is at least always a hop (office switches) away from the workers:

```

openqa:~ # salt -b 1 -l error -C 'G@roles:worker' cmd.run 'which iperf3 && iperf3 -c 2620:113:80c0:80a0:10:162:32:1f7'

```

Executing run on ['openqaworker2.suse.de']

```

jid:
  20201106124828227894
openqaworker2.suse.de:
  /usr/bin/iperf3
  Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
  [ 5] local 2620:113:80c0:8080:2e60:cff:fe73:2ac port 51558 connected to 2620:113:80c0:80a0:10:162:32:1f7
port 5201
  [ ID] Interval          Transfer          Bitrate          Retr  Cwnd

```

```

[ 5] 0.00-1.00 sec 102 MBytes 855 Mbits/sec 444 32.1 KBytes
[ 5] 1.00-2.00 sec 97.0 MBytes 814 Mbits/sec 338 18.1 KBytes
[ 5] 2.00-3.00 sec 95.6 MBytes 802 Mbits/sec 721 312 KBytes
[ 5] 3.00-4.00 sec 96.4 MBytes 809 Mbits/sec 628 34.9 KBytes
[ 5] 4.00-5.00 sec 92.2 MBytes 773 Mbits/sec 301 34.9 KBytes
[ 5] 5.00-6.00 sec 89.3 MBytes 749 Mbits/sec 494 113 KBytes
[ 5] 6.00-7.00 sec 87.3 MBytes 733 Mbits/sec 609 106 KBytes
[ 5] 7.00-8.00 sec 87.0 MBytes 730 Mbits/sec 325 251 KBytes
[ 5] 8.00-9.00 sec 86.3 MBytes 724 Mbits/sec 246 83.7 KBytes
[ 5] 9.00-10.00 sec 73.2 MBytes 614 Mbits/sec 93 142 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 906 MBytes 760 Mbits/sec 4199
[ 5] 0.00-10.04 sec 905 MBytes 756 Mbits/sec

```

iperf Done.

retcode:  
0

Executing run on ['powerqaworker-qam-1']

jid:

20201106124838673989

powerqaworker-qam-1:

/usr/bin/iperf3

Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201

[ 5] local 2620:113:80c0:80a0:10:162:30:de72 port 60628 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201

```

[ ID] Interval          Transfer      Bitrate      Retr  Cwnd
[ 5] 0.00-1.00 sec 104 MBytes 876 Mbits/sec 19 279 KBytes
[ 5] 1.00-2.00 sec 105 MBytes 881 Mbits/sec 2 286 KBytes
[ 5] 2.00-3.00 sec 105 MBytes 881 Mbits/sec 14 258 KBytes
[ 5] 3.00-4.00 sec 105 MBytes 881 Mbits/sec 6 255 KBytes
[ 5] 4.00-5.00 sec 106 MBytes 891 Mbits/sec 5 297 KBytes
[ 5] 5.00-6.00 sec 100 MBytes 839 Mbits/sec 8 252 KBytes
[ 5] 6.00-7.00 sec 108 MBytes 902 Mbits/sec 3 280 KBytes
[ 5] 7.00-8.00 sec 102 MBytes 860 Mbits/sec 5 322 KBytes
[ 5] 8.00-9.00 sec 108 MBytes 902 Mbits/sec 7 303 KBytes
[ 5] 9.00-10.00 sec 109 MBytes 912 Mbits/sec 10 261 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 1.03 GBytes 882 Mbits/sec 79
[ 5] 0.00-10.04 sec 1.02 GBytes 876 Mbits/sec

```

iperf Done.

retcode:  
0

Executing run on ['openqaworker13.suse.de']

jid:

20201106124849020759

openqaworker13.suse.de:

/usr/bin/iperf3

Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201

[ 5] local 2620:113:80c0:8080:10:160:2:26 port 53016 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201

```

[ ID] Interval          Transfer      Bitrate      Retr  Cwnd
[ 5] 0.00-1.00 sec 106 MBytes 893 Mbits/sec 27 247 KBytes
[ 5] 1.00-2.00 sec 107 MBytes 897 Mbits/sec 73 218 KBytes
[ 5] 2.00-3.00 sec 108 MBytes 902 Mbits/sec 83 204 KBytes
[ 5] 3.00-4.00 sec 103 MBytes 867 Mbits/sec 34 251 KBytes
[ 5] 4.00-5.00 sec 103 MBytes 866 Mbits/sec 55 132 KBytes
[ 5] 5.00-6.00 sec 105 MBytes 880 Mbits/sec 63 230 KBytes
[ 5] 6.00-7.00 sec 101 MBytes 849 Mbits/sec 129 201 KBytes
[ 5] 7.00-8.00 sec 104 MBytes 869 Mbits/sec 26 114 KBytes
[ 5] 8.00-9.00 sec 104 MBytes 873 Mbits/sec 60 218 KBytes
[ 5] 9.00-10.00 sec 103 MBytes 865 Mbits/sec 51 211 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 1.02 GBytes 876 Mbits/sec 601
[ 5] 0.00-10.04 sec 1.02 GBytes 870 Mbits/sec

```

iperf Done.

retcode:

0

Executing run on ['openqaworker10.suse.de']

```

jid:
  20201106124859469998
openqaworker10.suse.de:
  /usr/bin/iperf3
  Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
  [ 5] local 2620:113:80c0:8080:10:160:68:1 port 38128 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5
201
  [ ID] Interval          Transfer      Bitrate      Retr  Cwnd
  [ 5]  0.00-1.00    sec    111 MBytes    930 Mb/s     51   208 KBytes
  [ 5]  1.00-2.00    sec    107 MBytes    900 Mb/s    132   199 KBytes
  [ 5]  2.00-3.00    sec    106 MBytes    892 Mb/s    126  69.7 KBytes
  [ 5]  3.00-4.00    sec    106 MBytes    891 Mb/s    115   159 KBytes
  [ 5]  4.00-5.00    sec    105 MBytes    879 Mb/s    125   279 KBytes
  [ 5]  5.00-6.00    sec    109 MBytes    911 Mb/s     75   252 KBytes
  [ 5]  6.00-7.00    sec    104 MBytes    869 Mb/s    124   291 KBytes
  [ 5]  7.00-8.00    sec    107 MBytes    894 Mb/s    130   216 KBytes
  [ 5]  8.00-9.00    sec    109 MBytes    915 Mb/s     42   199 KBytes
  [ 5]  9.00-10.00   sec    107 MBytes    898 Mb/s    128   223 KBytes
  -----
  [ ID] Interval          Transfer      Bitrate      Retr
  [ 5]  0.00-10.00   sec    1.05 GBytes    898 Mb/s    1048
  [ 5]  0.00-10.05   sec    1.04 GBytes    892 Mb/s
  sender
  receiver

```

iperf Done.

retcode:
0

Executing run on ['QA-Power8-5-kvm.qa.suse.de']

```

QA-Power8-5-kvm.qa.suse.de:
  /usr/bin/iperf3
  Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
  [ 5] local 2620:113:80c0:80a0:10:162:2a:5c8d port 52542 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201
t 5201
  [ ID] Interval          Transfer      Bitrate      Retr  Cwnd
  [ 5]  0.00-1.00    sec    89.5 MBytes    751 Mb/s    327   325 KBytes
  [ 5]  1.00-2.00    sec    92.0 MBytes    772 Mb/s    624   250 KBytes
  [ 5]  2.00-3.00    sec    98.5 MBytes    826 Mb/s    490  73.9 KBytes
  [ 5]  3.00-4.00    sec    94.6 MBytes    793 Mb/s    607   152 KBytes
  [ 5]  4.00-5.00    sec    96.2 MBytes    807 Mb/s    521   445 KBytes
  [ 5]  5.00-6.00    sec    95.7 MBytes    803 Mb/s    833   34.9 KBytes
  [ 5]  6.00-7.01    sec    95.8 MBytes    799 Mb/s    787   78.1 KBytes
  [ 5]  7.01-8.00    sec    89.4 MBytes    755 Mb/s    980   181 KBytes
  [ 5]  8.00-9.00    sec    91.4 MBytes    767 Mb/s    243   137 KBytes
  [ 5]  9.00-10.00   sec    73.5 MBytes    616 Mb/s    515   25.1 KBytes
  -----
  [ ID] Interval          Transfer      Bitrate      Retr
  [ 5]  0.00-10.00   sec    917 MBytes    769 Mb/s    5927
  [ 5]  0.00-10.04   sec    914 MBytes    764 Mb/s
  sender
  receiver

```

iperf Done.

jid:
20201106124909926319
retcode:
0

Executing run on ['openqaworker5.suse.de']

```

jid:
  20201106124920344427
openqaworker5.suse.de:
  /usr/bin/iperf3
  Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
  [ 5] local 2620:113:80c0:8080:10:160:1:93 port 50440 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5
201
  [ ID] Interval          Transfer      Bitrate      Retr  Cwnd
  [ 5]  0.00-1.00    sec    105 MBytes    877 Mb/s    309   107 KBytes
  [ 5]  1.00-2.00    sec    99.8 MBytes    837 Mb/s    337   110 KBytes
  [ 5]  2.00-3.00    sec    103 MBytes    868 Mb/s    154   188 KBytes
  [ 5]  3.00-4.00    sec    99.8 MBytes    837 Mb/s    377   314 KBytes
  [ 5]  4.00-5.00    sec    100 MBytes    843 Mb/s    432   86.5 KBytes

```

```

[ 5] 5.00-6.00 sec 99.5 MBytes 835 Mb/s 310 234 KBytes
[ 5] 6.00-7.00 sec 104 MBytes 872 Mb/s 222 206 KBytes
[ 5] 7.00-8.00 sec 99.5 MBytes 834 Mb/s 246 107 KBytes
[ 5] 8.00-9.00 sec 98.8 MBytes 829 Mb/s 290 251 KBytes
[ 5] 9.00-10.00 sec 96.6 MBytes 811 Mb/s 465 155 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 1006 MBytes 844 Mb/s 3142
[ 5] 0.00-10.04 sec 1004 MBytes 839 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['openqaworker8.suse.de']

jid:

20201106124930709117

openqaworker8.suse.de:

/usr/bin/iperf3

Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201

```
[ 5] local 2620:113:80c0:8080:ec4:7aff:fe99:dc5b port 54914 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201
```

```

[ ID] Interval          Transfer      Bitrate      Retr  Cwnd
[ 5] 0.00-1.00 sec 96.2 MBytes 807 Mb/s 824 26.5 KBytes
[ 5] 1.00-2.00 sec 94.3 MBytes 791 Mb/s 404 160 KBytes
[ 5] 2.00-3.00 sec 87.8 MBytes 737 Mb/s 510 26.5 KBytes
[ 5] 3.00-4.00 sec 95.4 MBytes 800 Mb/s 709 230 KBytes
[ 5] 4.00-5.00 sec 98.5 MBytes 827 Mb/s 604 127 KBytes
[ 5] 5.00-6.00 sec 93.0 MBytes 780 Mb/s 709 32.1 KBytes
[ 5] 6.00-7.00 sec 97.8 MBytes 820 Mb/s 419 75.3 KBytes
[ 5] 7.00-8.00 sec 94.6 MBytes 793 Mb/s 605 93.4 KBytes
[ 5] 8.00-9.00 sec 102 MBytes 853 Mb/s 484 244 KBytes
[ 5] 9.00-10.00 sec 46.6 MBytes 391 Mb/s 78 60.0 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 906 MBytes 760 Mb/s 5346
[ 5] 0.00-10.04 sec 904 MBytes 755 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['openqaworker9.suse.de']

jid:

20201106124941047009

openqaworker9.suse.de:

/usr/bin/iperf3

Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201

```
[ 5] local 2620:113:80c0:8080:10:160:1:20 port 34090 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201
```

```

[ ID] Interval          Transfer      Bitrate      Retr  Cwnd
[ 5] 0.00-1.00 sec 57.0 MBytes 478 Mb/s 3 713 KBytes
[ 5] 1.00-2.00 sec 55.0 MBytes 461 Mb/s 90 307 KBytes
[ 5] 2.00-3.00 sec 101 MBytes 849 Mb/s 109 71.1 KBytes
[ 5] 3.00-4.00 sec 100 MBytes 839 Mb/s 522 293 KBytes
[ 5] 4.00-5.00 sec 102 MBytes 860 Mb/s 212 211 KBytes
[ 5] 5.00-6.00 sec 101 MBytes 849 Mb/s 342 269 KBytes
[ 5] 6.00-7.00 sec 86.2 MBytes 724 Mb/s 499 276 KBytes
[ 5] 7.00-8.00 sec 48.8 MBytes 409 Mb/s 1401 37.7 KBytes
[ 5] 8.00-9.00 sec 71.2 MBytes 598 Mb/s 576 170 KBytes
[ 5] 9.00-10.00 sec 96.2 MBytes 807 Mb/s 640 218 KBytes

```

```

-----
[ ID] Interval          Transfer      Bitrate      Retr
[ 5] 0.00-10.00 sec 820 MBytes 687 Mb/s 4394
[ 5] 0.00-10.04 sec 816 MBytes 682 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['QA-Power8-4-kvm.qa.suse.de']

QA-Power8-4-kvm.qa.suse.de:

```

/usr/bin/iperf3
Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
[ 5] local 2620:113:80c0:80a0:10:162:31:3446 port 53754 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201

```

[ ID]	Interval		Transfer	Bitrate	Retr	Cwnd	
[ 5]	0.00-1.00	sec	106 MBytes	889 Mb/s	17	180 KBytes	
[ 5]	1.00-2.00	sec	101 MBytes	843 Mb/s	1	377 KBytes	
[ 5]	2.00-3.00	sec	100 MBytes	842 Mb/s	35	198 KBytes	
[ 5]	3.00-4.00	sec	104 MBytes	871 Mb/s	19	78.1 KBytes	
[ 5]	4.00-5.00	sec	102 MBytes	859 Mb/s	18	322 KBytes	
[ 5]	5.00-6.00	sec	84.8 MBytes	711 Mb/s	2	282 KBytes	
[ 5]	6.00-7.00	sec	89.4 MBytes	750 Mb/s	14	257 KBytes	
[ 5]	7.00-8.00	sec	103 MBytes	860 Mb/s	8	279 KBytes	
[ 5]	8.00-9.00	sec	100 MBytes	843 Mb/s	9	298 KBytes	
[ 5]	9.00-10.00	sec	104 MBytes	876 Mb/s	6	245 KBytes	

---

```

[ ID] Interval          Transfer      Bitrate      Retr
[ 5]  0.00-10.00 sec    995 MBytes  834 Mb/s    129
[ 5]  0.00-10.06 sec    992 MBytes  828 Mb/s

```

```

iperf Done.
jid: 20201106124951627941
retcode: 0

```

Executing run on ['openqaworker3.suse.de']

```

jid: 20201106125002001609
openqaworker3.suse.de:
  which: no iperf3 in (/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin)
retcode: 1

```

Executing run on ['grenache-1.qa.suse.de']

```

grenache-1.qa.suse.de:
/usr/bin/iperf3
Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
[ 5] local 2620:113:80c0:80a0:10:162:29:12f0 port 43282 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201

```

[ ID]	Interval		Transfer	Bitrate	Retr	Cwnd	
[ 5]	0.00-1.00	sec	113 MBytes	952 Mb/s	15	169 KBytes	
[ 5]	1.00-2.00	sec	110 MBytes	923 Mb/s	3	176 KBytes	
[ 5]	2.00-3.00	sec	110 MBytes	923 Mb/s	10	294 KBytes	
[ 5]	3.00-4.00	sec	110 MBytes	923 Mb/s	11	170 KBytes	
[ 5]	4.00-5.00	sec	111 MBytes	933 Mb/s	11	329 KBytes	
[ 5]	5.00-6.00	sec	109 MBytes	912 Mb/s	12	280 KBytes	
[ 5]	6.00-7.00	sec	111 MBytes	933 Mb/s	15	153 KBytes	
[ 5]	7.00-8.00	sec	110 MBytes	923 Mb/s	11	315 KBytes	
[ 5]	8.00-9.00	sec	110 MBytes	923 Mb/s	11	163 KBytes	
[ 5]	9.00-10.00	sec	110 MBytes	923 Mb/s	13	149 KBytes	

---

```

[ ID] Interval          Transfer      Bitrate      Retr
[ 5]  0.00-10.00 sec    1.08 GBytes  927 Mb/s    112
[ 5]  0.00-10.05 sec    1.08 GBytes  920 Mb/s

```

```

iperf Done.
jid: 20201106125002233677
retcode: 0

```

Executing run on ['openqaworker-arm-2.suse.de']

```

jid: 20201106125012911475
openqaworker-arm-2.suse.de:
/usr/bin/iperf3
Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
[ 5] local 2620:113:80c0:8080:1e1b:dff:fe68:ee4d port 48828 connected to 2620:113:80c0:80a0:10:162:32:1f7 port 5201

```

[ ID]	Interval		Transfer	Bitrate	Retr	Cwnd	
[ 5]	0.00-1.00	sec	112 MBytes	936 Mb/s	0	3.00 MBytes	



```

[ 5] 1.00-2.00 sec 109 MBytes 912 Mb/s 0 3.00 MBytes
[ 5] 2.00-3.00 sec 109 MBytes 912 Mb/s 0 3.00 MBytes
[ 5] 3.00-4.00 sec 109 MBytes 912 Mb/s 0 3.00 MBytes
[ 5] 4.00-5.00 sec 106 MBytes 892 Mb/s 0 3.00 MBytes
[ 5] 5.00-6.00 sec 98.8 MBytes 828 Mb/s 0 3.00 MBytes
[ 5] 6.00-7.00 sec 110 MBytes 923 Mb/s 0 3.00 MBytes
[ 5] 7.00-8.00 sec 110 MBytes 923 Mb/s 0 3.00 MBytes
[ 5] 8.00-9.00 sec 111 MBytes 933 Mb/s 0 3.15 MBytes
[ 5] 9.00-10.00 sec 110 MBytes 923 Mb/s 0 3.15 MBytes

```

```

-----
[ ID] Interval      Transfer    Bitrate    Retr
[ 5] 0.00-10.00 sec 1.06 GBytes 909 Mb/s   0
[ 5] 0.00-10.03 sec 1.06 GBytes 907 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['openqaworker-arm-1.suse.de']

```

jid:
20201106125023654762
openqaworker-arm-1.suse.de:
/usr/bin/iperf3
Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
[ 5] local 2620:113:80c0:8080:1e1b:dff:fe68:7ec7 port 60232 connected to 2620:113:80c0:80a0:10:162:32:1f7
port 5201

```

```

[ ID] Interval      Transfer    Bitrate    Retr  Cwnd
[ 5] 0.00-1.00 sec 114 MBytes 954 Mb/s   0 3.00 MBytes
[ 5] 1.00-2.00 sec 106 MBytes 892 Mb/s   0 3.00 MBytes
[ 5] 2.00-3.00 sec 108 MBytes 902 Mb/s   0 3.00 MBytes
[ 5] 3.00-4.00 sec 108 MBytes 902 Mb/s   0 3.00 MBytes
[ 5] 4.00-5.00 sec 109 MBytes 912 Mb/s   0 3.00 MBytes
[ 5] 5.00-6.00 sec 110 MBytes 923 Mb/s   0 3.00 MBytes
[ 5] 6.00-7.00 sec 109 MBytes 912 Mb/s   0 3.00 MBytes
[ 5] 7.00-8.00 sec 109 MBytes 912 Mb/s   0 3.00 MBytes
[ 5] 8.00-9.00 sec 109 MBytes 912 Mb/s   0 3.00 MBytes
[ 5] 9.00-10.00 sec 110 MBytes 923 Mb/s   0 3.00 MBytes

```

```

-----
[ ID] Interval      Transfer    Bitrate    Retr
[ 5] 0.00-10.00 sec 1.06 GBytes 914 Mb/s   0
[ 5] 0.00-10.03 sec 1.06 GBytes 912 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['malbec.arch.suse.de']

```

jid:
20201106125034314896
malbec.arch.suse.de:
/usr/bin/iperf3
Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
[ 5] local 2620:113:80c0:8000:10:161:24:54 port 38826 connected to 2620:113:80c0:80a0:10:162:32:1f7 port
5201

```

```

[ ID] Interval      Transfer    Bitrate    Retr  Cwnd
[ 5] 0.00-1.00 sec 107 MBytes 894 Mb/s   7 148 KBytes
[ 5] 1.00-2.00 sec 105 MBytes 881 Mb/s   2 381 KBytes
[ 5] 2.00-3.00 sec 102 MBytes 860 Mb/s   3 259 KBytes
[ 5] 3.00-4.00 sec 104 MBytes 870 Mb/s   4 291 KBytes
[ 5] 4.00-5.00 sec 102 MBytes 860 Mb/s   5 322 KBytes
[ 5] 5.00-6.00 sec 108 MBytes 902 Mb/s   7 305 KBytes
[ 5] 6.00-7.00 sec 105 MBytes 881 Mb/s   2 296 KBytes
[ 5] 7.00-8.00 sec 105 MBytes 881 Mb/s   7 276 KBytes
[ 5] 8.00-9.00 sec 102 MBytes 860 Mb/s   6 284 KBytes
[ 5] 9.00-10.00 sec 101 MBytes 849 Mb/s   3 317 KBytes

```

```

-----
[ ID] Interval      Transfer    Bitrate    Retr
[ 5] 0.00-10.00 sec 1.02 GBytes 874 Mb/s  46
[ 5] 0.00-10.05 sec 1.01 GBytes 867 Mb/s

```

iperf Done.

```
retcode:
0
```

Executing run on ['openqaworker6.suse.de']

```
jid:
  20201106125044714148
openqaworker6.suse.de:
  /usr/bin/iperf3
  Connecting to host 2620:113:80c0:80a0:10:162:32:1f7, port 5201
  [ 5] local 2620:113:80c0:8080:10:160:1:100 port 45096 connected to 2620:113:80c0:80a0:10:162:32:1f7 port
5201
  [ ID] Interval           Transfer     Bitrate      Retr  Cwnd
  [ 5]  0.00-1.00    sec     101 MBytes    847 Mbits/sec  252   141 KBytes
  [ 5]  1.00-2.00    sec     101 MBytes    843 Mbits/sec  275   209 KBytes
  [ 5]  2.00-3.00    sec    99.3 MBytes    833 Mbits/sec  349   99.0 KBytes
  [ 5]  3.00-4.00    sec    96.2 MBytes    807 Mbits/sec  314   243 KBytes
  [ 5]  4.00-5.00    sec    100 MBytes    841 Mbits/sec  424   144 KBytes
  [ 5]  5.00-6.00    sec    79.6 MBytes    668 Mbits/sec  284   250 KBytes
  [ 5]  6.00-7.00    sec    98.8 MBytes    829 Mbits/sec  327   93.4 KBytes
  [ 5]  7.00-8.00    sec     101 MBytes    848 Mbits/sec  336   145 KBytes
  [ 5]  8.00-9.00    sec    97.7 MBytes    820 Mbits/sec  468   106 KBytes
  [ 5]  9.00-10.00   sec    97.6 MBytes    818 Mbits/sec  345   144 KBytes
  - - - - -
  [ ID] Interval           Transfer     Bitrate      Retr
  [ 5]  0.00-10.00   sec     972 MBytes    815 Mbits/sec  3374
  [ 5]  0.00-10.04   sec     970 MBytes    810 Mbits/sec
                                     sender
                                     receiver

  iperf Done.
retcode:
  0
```

So with these numbers I'm pretty certain that everything works as expected.

#### #20 - 2020-11-09 16:33 - cdywan

- Status changed from In Progress to Resolved

nicksinger wrote:

After applying these changes, OSD can be reached over v6 from all machines:

[...]

So with these numbers I'm pretty certain that everything works as expected.

So the ticket is *Resolved* I take it?

#### #21 - 2020-11-09 18:45 - okurz

- Status changed from Resolved to In Progress

As this ticket was about an issue causing lots of problems and confusions but was also caused by the team itself I would really keep it open and up for the assignee to decide when it is "Resolved". Definitely I think an issue specific retrospective should be conducted

Also

<https://infra.nue.suse.com/SelfService/Display.html?id=178626>

is still open

#### #22 - 2020-11-10 17:29 - cdywan

- Due date changed from 2020-10-24 to 2020-11-13

Ack

#### #23 - 2020-11-11 07:43 - nicksinger

besides what was mentioned by Oli we also need a proper permanent solution in salt

#### #24 - 2020-11-11 07:56 - okurz

As you wrote we need to set `net.ipv6.conf.$main_interface.accept_ra = 2`

To get `$main_interface` <https://tedops.github.io/how-to-find-default-active-ethernet-interface.html> looks promising, e.g. call

```
salt '*' network.default_route inet
```

I guess in salt state files we should do:

```
net.ipv6.conf.{{ salt['network.default_route']('inet')[0]['interface'] }}.accept_ra:
  sysctl.present:
    - value: 2
```

if this does not work then probably a custom grain function should be used, as in <https://lemarchand.io/saltstack-and-internal-network-interfaces/>

#### #25 - 2020-11-13 11:54 - nicksinger

- Description updated

#### #26 - 2020-11-13 11:56 - nicksinger

- Description updated

#### #27 - 2020-11-16 11:50 - nicksinger

- Has duplicate action #77995: worker instances on grenache-1 seem to fail (sometimes?) to connect to web-uis added

#### #28 - 2020-11-17 11:07 - okurz

Trying the suggestion from [#73633#note-24](#) on osd with a temporary change to /srv/salt/openqa/worker.sls and trying to apply with salt 'openqaworker10\*' state.apply test=True I get:

```
openqaworker10.suse.de:
  Data failed to compile:
  -----
  Rendering SLS 'base:openqa.worker' failed: Jinja error: 'anycast' does not appear to be an IPv4 or IPv6 network
Traceback (most recent call last):
  File "/usr/lib/python3.6/site-packages/salt/utils/templates.py", line 394, in render_jinja_tmpl
    output = template.render(**decoded_context)
  File "/usr/lib/python3.6/site-packages/jinja2/asyncsupport.py", line 76, in render
    return original_render(self, *args, **kwargs)
  File "/usr/lib/python3.6/site-packages/jinja2/environment.py", line 1008, in render
    return self.environment.handle_exception(exc_info, True)
  File "/usr/lib/python3.6/site-packages/jinja2/environment.py", line 780, in handle_exception
    reraise(exc_type, exc_value, tb)
  File "/usr/lib/python3.6/site-packages/jinja2/_compat.py", line 37, in reraise
    raise value.with_traceback(tb)
  File "<template>", line 367, in top-level template code
  File "/usr/lib/python3.6/site-packages/salt/modules/network.py", line 1690, in default_route
    _routes = routes()
  File "/usr/lib/python3.6/site-packages/salt/modules/network.py", line 1647, in routes
    routes_ = _ip_route_linux()
  File "/usr/lib/python3.6/site-packages/salt/modules/network.py", line 569, in _ip_route_linux
    address_mask = convert_cidr(comps[0])
  File "/usr/lib/python3.6/site-packages/salt/modules/network.py", line 1149, in convert_cidr
    cidr = calc_net(cidr)
  File "/usr/lib/python3.6/site-packages/salt/modules/network.py", line 1171, in calc_net
    return salt.utils.network.calc_net(ip_addr, netmask)
  File "/usr/lib/python3.6/site-packages/salt/utils/network.py", line 1053, in calc_net
    return six.text_type(ipaddress.ip_network(ip_addr, strict=False))
  File "/usr/lib64/python3.6/ipaddress.py", line 84, in ip_network
    address)
ValueError: 'anycast' does not appear to be an IPv4 or IPv6 network

; line 367
```

#### #29 - 2020-11-17 11:23 - okurz

- Related to action #68095: Migrate osd workers from SuSEfirewall2 to firewalld added

#### #30 - 2020-11-17 12:44 - okurz

Trying a "5 Whys" analysis.

First mkittler worked on migrating SuSEfirewall2 to firewalld in [#68095](#) . On 2020-10-19 13:20 CEST the according salt change was deployed to all workers

We were informed about a "general problem" by our monitoring and also by user reports about 2h later. Even before 2020-10-20 12:46 CEST nicksinger has manually added routes to workers as described in [#75055](#) which then caused further issues. This looked good because as [#75055](#) states: "the worker appeared on all webui's again" but the performance decreased heavily and lead to [#73633](#) .

Maybe there

- Why did we not see any problems directly after the salt state was applied?
  - it was not "completely broken" and took 24h to trigger the big alert, likely just after nsinger applied additional changes
  - -> suggestion: We should have monitoring for basic "IPv4 and IPv6 works to ping, top, http from all machines to all machines". Make sure to *explicitly* select both stacks
  - -> suggestion: A passive performance measurement regarding throughput on interfaces
- Why did we not already have a ticket for the issue that mmoese reported on 2020-10-20?
  - At the time we did not see "baremetal-support.qa.suse.de" as that important for us and could not link it to an issue in the general osd infrastructure.
  - -> suggestion: whenever we apply changes to the infrastructure we should have a ticket
  - TODO lookup the according infra ticket and check when it was created
  - -> suggestion: Whenever creating any external ticket, e.g. Englnfra, create internal tracker ticket. Because there might be more internal notes
- Why did we not see the connection the firewalld migration [#68095](#) ?
  - Because no tests directly linked to the ticket or deployed salt changes failed
  - -> suggestion: Same as in OSD deployment we should look for failed grafana
  - -> suggestion: Collect all the information between "last good" and "first bad" and then also find the git diff in openqa/salt-states-openqa
- Why did mkittler and me think that the firewalld change was not the issue?
  - We thought firewalld was "long gone" because mkittler already created the SR at 2020-10-15 (but only partially deployed for better testing)
  - We jumped to the conclusion that IPv6 changes within the network out of our control should have triggered that
  - -> suggestion: Apply proper "scientific method" with written down hypotheses, experiments and conclusions in tickets, follow <https://progress.opensuse.org/projects/openqav3/wiki#Further-decision-steps-working-on-test-issues>
  - -> suggestion: Keep salt states to describe what should *not* be there
  - -> suggestion: Try out older brifs snapshots in systems for crosschecking and boot with disabled salt. In the kernel cmdline append `systemd.mask=salt-minion.service`
- Why did it take so long?
  - Because Englnfra was too slow to tell us it's not their fault
  - nicksinger did not get an answer for "long enough" so he figured it's our own fault
  - We thought "good enough workarounds are in place" and worked on other tickets that helped to resolve the actual issue, e.g. [#75055](#) , [#75016](#)
  - -> Conclusion: Actually we did good because the user base was not impacted that much anymore, we had workarounds in place, we were investigating other issues but always kept the relation to this ticket in mind which in the end helped to fix it
- Why are we still not finished?
  - Because cdywan does not run dailies to check on urgent tickets still open
  - -> suggestion: team should conduct a work backlog check on a daily base
  - We were not sure if any other person should take the ticket from nsinger
  - -> suggestion: nsinger does not mind if someone else provides a suggestion or takes over the ticket

-> [#78127](#)

#### **#31 - 2020-11-17 12:48 - okurz**

- Copied to action [#78127](#): follow-up to [#73633](#) - lessons learned and suggestions added

#### **#32 - 2020-11-17 19:14 - cdywan**

- Due date changed from 2020-11-13 to 2020-11-17

#### **#33 - 2020-11-18 09:34 - nicksinger**

I'll take over from here and will try to implement a proper salt solution. This is my plan of action :)

#### **#34 - 2020-11-18 10:52 - nicksinger**

- Status changed from In Progress to Feedback

[https://gitlab.suse.de/openqa/salt-states-openqa/-/merge\\_requests/401](https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/401) - anyone disagrees that we can close this here once this MR is merged? :)

### #35 - 2020-11-18 13:55 - okurz

well, the deploy pipeline failed so I suggest to resolve this ticket as soon as you can check that this setting is actually applied on all affected machines :)

And are there still workarounds in place that we need to remove?

### #36 - 2020-11-19 15:01 - cdywan

okurz wrote:

well, the deploy pipeline failed so I suggest to resolve this ticket as soon as you can check that this setting is actually applied on all affected machines :)

And are there still workarounds in place that we need to remove?

I re-ran the [pipeline](#) on *master* and *deploy* failed like this:

```
RROR: Minions returned with non-zero exit code
1411openqaworker-arm-1.suse.de:
1412Summary for openqaworker-arm-1.suse.de
1413-----
1414Succeeded: 285
1415Failed: 0
```

### #37 - 2020-11-19 20:12 - okurz

If in the gitlab CI pipeline job log you scroll further to the top you can find <https://gitlab.suse.de/openqa/salt-states-openqa/-/jobs/289706#L822> which says:

```
openqa-monitor.qa.suse.de:
  Data failed to compile:
-----
  Rendering SLS 'base:openqa.monitoring.grafana' failed: while constructing a mapping
  in "<unicode string>", line 10, column 1
  found conflicting ID '/var/lib/grafana/dashboards//worker-localhost.json'
  in "<unicode string>", line 193, column 1
openqaworker2.suse.de:
```

which we have a ticket about: [#75445](#) which seems to be causing more problems now, hence raising prio there.

### #38 - 2020-11-20 09:18 - okurz

looked into the topic together with nsinger:  
Experimented on openqaworker10:

```
# dig openqa.suse.de AAAA
dig: parse of /etc/resolv.conf failed
```

looking into `/etc/resolv.conf` which was from 2018-10-19 with content:

```
search suse.de
nameserver fe80::20d:b9ff:fe01:ea8%eth2
nameserver 10.160.0.1
```

calling `netconfig update -f` replaced the file with a symlink. I remember that there was some system upgrade where one should have replaced manually maintained files with symlinks like this. Probably a good idea to do that on all our machines.

Now we can test again properly:

```
# dig openqa.suse.de AAAA

; <<>> DiG 9.16.6 <<>> openqa.suse.de AAAA
;; global options: +cmd
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 12530
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 7

;; OPT PSEUDOSECTION:
;; EDNS: version: 0, flags:; udp: 4096
;; COOKIE: 91ba81fa158dab233fb0c3735fb788ad2f77d9cbe5e72c49 (good)
;; QUESTION SECTION:
;openqa.suse.de. IN AAAA

;; ANSWER SECTION:
```

```
openqa.suse.de.      300 IN AAAA      2620:113:80c0:8080:10:160:0:207
```

```
;; AUTHORITY SECTION:
```

```
suse.de.      300 IN NS      dns1.suse.de.  
suse.de.      300 IN NS      frao-p-infoblox-01.corp.suse.com.  
suse.de.      300 IN NS      dns2.suse.de.  
suse.de.      300 IN NS      frao-p-infoblox-02.corp.suse.com.
```

```
;; ADDITIONAL SECTION:
```

```
dns2.suse.de.      300 IN AAAA      2620:113:80c0:8080:10:160:0:1  
dns1.suse.de.      300 IN AAAA      2620:113:80c0:8080:10:160:2:88  
dns2.suse.de.      300 IN A      10.160.0.1  
dns1.suse.de.      300 IN A      10.160.2.88  
frao-p-infoblox-02.corp.suse.com. 14863 IN A      10.156.86.70  
frao-p-infoblox-01.corp.suse.com. 14863 IN A      10.156.86.6
```

```
;; Query time: 0 msec
```

```
;; SERVER: 2620:113:80c0:8080:10:160:0:1#53(2620:113:80c0:8080:10:160:0:1)
```

```
;; WHEN: Fri Nov 20 10:13:17 CET 2020
```

```
;; MSG SIZE rcvd: 336
```

```
# which iperf3 && iperf3 -c 2620:113:80c0:8080:10:160:0:207
```

```
/usr/bin/iperf3
```

```
Connecting to host 2620:113:80c0:8080:10:160:0:207, port 5201
```

```
[ 5] local 2620:113:80c0:8080:10:160:68:35 port 43226 connected to 2620:113:80c0:8080:10:160:0:207 port 5201
```

[ ID]	Interval	Transfer	Bitrate	Retr	Cwnd	
[ 5]	0.00-1.00	sec 110 MBytes	927 Mbits/sec	9	223 KBytes	
[ 5]	1.00-2.00	sec 108 MBytes	910 Mbits/sec	8	213 KBytes	
[ 5]	2.00-3.00	sec 109 MBytes	916 Mbits/sec	2	298 KBytes	
[ 5]	3.00-4.00	sec 108 MBytes	909 Mbits/sec	4	286 KBytes	
[ 5]	4.00-5.00	sec 110 MBytes	925 Mbits/sec	9	205 KBytes	
[ 5]	5.00-6.00	sec 107 MBytes	894 Mbits/sec	9	220 KBytes	
[ 5]	6.00-7.00	sec 109 MBytes	915 Mbits/sec	4	159 KBytes	
[ 5]	7.00-8.00	sec 110 MBytes	919 Mbits/sec	8	149 KBytes	
[ 5]	8.00-9.00	sec 108 MBytes	904 Mbits/sec	5	259 KBytes	
[ 5]	9.00-10.00	sec 107 MBytes	900 Mbits/sec	5	216 KBytes	
-----						
[ ID]	Interval	Transfer	Bitrate	Retr		
[ 5]	0.00-10.00	sec 1.06 GBytes	912 Mbits/sec	63		sender
[ 5]	0.00-10.03	sec 1.06 GBytes	907 Mbits/sec			receiver

```
iperf Done.
```

same for iperf3 -6 -c openqa.suse.de. So this looks good so far, same should be applied to all machines, simply over salt seems safe. Done that.

### #39 - 2020-11-20 10:41 - nicksinger

I've brought back the two power workers (malbec and powerqaworker-qam-1). I see ping fails on the following workers: openqaworker8.suse.de, openqaworker-arm-1.suse.de and openqaworker-arm-2.suse.de which is expected:

```
openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'ls -lah /etc/sysctl.d/poo73633_debugging.conf && echo "workaround in place" || true'  
openqaworker2.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
openqaworker3.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
openqaworker6.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
openqaworker5.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
openqaworker8.suse.de:  
-rw-r--r-- 1 root root 35 Oct 24 13:27 /etc/sysctl.d/poo73633_debugging.conf  
workaround in place  
openqaworker9.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
powerqaworker-qam-1:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
QA-Power8-4-kvm.qa.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
QA-Power8-5-kvm.qa.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
malbec.arch.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory  
openqaworker10.suse.de:  
ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory
```

```
openqaworker13.suse.de:
  ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory
grenache-1.qa.suse.de:
  ls: cannot access '/etc/sysctl.d/poo73633_debugging.conf': No such file or directory
openqaworker-arm-1.suse.de:
  -rw-r--r-- 1 root root 35 Oct 22 19:29 /etc/sysctl.d/poo73633_debugging.conf
  workaround in place
openqaworker-arm-2.suse.de:
  -rw-r--r-- 1 root root 35 Oct 22 19:30 /etc/sysctl.d/poo73633_debugging.conf
  workaround in place
```

I removed these files now and changed the running value with `openqa:~ # salt -l error -C 'G@roles:worker' cmd.run 'sysctl net.ipv6.conf.all.disable_ipv6=0'`. I reran the iperf-check and saw >800MB/s for all hosts. The salt change is persisted (in /etc/sysctl.d/99-salt.conf) and also the runtime configuration is set to `net.ipv6.conf.$default_interface.accept_ra = 2`. I would consider this now as finally done. Any objections? :)

**#40 - 2020-11-20 11:46 - okurz**

- Status changed from Feedback to Resolved

Thanks, perfect final actions :)

**#41 - 2020-11-21 07:27 - okurz**

- Related to action #80128: openqaworker-arm-2 fails to download from openqa added

**Files**

---

ip6tables-save.firewalld.txt	5.98 KB	2020-11-04	nicksinger
ip6tables-save.susefirewall.txt	3.73 KB	2020-11-04	nicksinger