

openQA Project - action #35895

[os-autoinst-openvswitch][theory] problem with assigning VLANs tags - more testsuites are using the same VLAN?

2018-05-04 10:18 - thehejik

Status:	Resolved	Start date:	2018-05-04
Priority:	Normal	Due date:	
Assignee:	EDiGiacinto	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Done		
Difficulty:			
Description			
Anton noticed an issue when support_server failed due network problems - ss leased 10.0.2.22 most probably from different ss in the same VLAN.			
https://openqa.suse.de/tests/1667461 failed support_server oqw3 tap4 3.5. 20:33(+2 on worker)			
https://openqa.suse.de/tests/1667530 parallel_failed master oqw7 tap12 3.5. 20:33			
https://openqa.suse.de/tests/1667523 parallel_failed slave oqw6 tap0 3.5. 20:33			
So I did a check on oqw{3,7,6} and found a potential issue that on oqw7 we have more than one tap device assigned to VLAN tag=50.			
Anton has added OVS_DEBUG=1 to his testsuites so with another occurrence of this issue we can debug it in autoinst-log.txt (and see if some foreign taps are assigned to the same VLAN).			
note to myself ... determine VLAN tag, see test start time => sudo journalctl -u os-autoinst-openvswitch grep -e "tag.*50"			
thehejik@openqaworker7:~> sudo journalctl -u os-autoinst-openvswitch grep -e "tag.*50"			
kvě 03 22:29:50 openqaworker7 ovs-vsctl[16033]: ovs 00001 vsctl INFO Called as ovs-vsctl set port tap12 tag=50			
kvě 03 22:29:51 openqaworker7 ovs-vsctl[16058]: ovs 00001 vsctl INFO Called as ovs-vsctl set port tap13 tag=50			
kvě 03 22:33:18 openqaworker7 ovs-vsctl[16899]: ovs 00001 vsctl INFO Called as ovs-vsctl set port tap12 tag=50			
kvě 03 22:51:39 openqaworker7 ovs-vsctl[24378]: ovs 00001 vsctl INFO Called as ovs-vsctl set port tap9 tag=50			
kvě 03 22:52:32 openqaworker7 ovs-vsctl[24896]: ovs 00001 vsctl INFO Called as ovs-vsctl remove port tap9 tag 50			
Related problem can be that VLAN tags are probably not removed when the test fail.			

History

#1 - 2018-05-28 06:15 - okurz

This is an autogenerated message for openQA integration by the openqa_review script:

This bug is still referenced in a failing openQA test: hpc_ganglia_supportserver

<https://openqa.suse.de/tests/1728938>

#2 - 2018-06-22 08:32 - mlovicka

Similar issue happen in following run <https://openqa.suse.de/tests/177426>.

output from serial0.txt:

```
Jun 21 15:36:40 susetest systemd[1]: Starting wicked managed network interfaces...
Jun 21 15:36:40 susetest wickedd[964]: eth0: IP address 10.0.2.1 is in use by 52:54:00:12:03:9e
Jun 21 15:36:40 susetest wickedd[964]: eth0: address '10.0.2.1' is already in use
Jun 21 15:36:40 susetest wickedd[964]: __ni_rtnl_send_newroute(ipv4 0.0.0.0/0 via 10.0.2.2 dev eth0 type unicast table main scope universe protocol boot): ni_nl_talk failed [Unspecific failure]
Jun 21 15:37:10 susetest wickedd[1718]: lo up
Jun 21 15:37:10 susetest wickedd[1718]: eth0 device-not-running
Jun 21 15:37:10 susetest systemd[1]: Started wicked managed network interfaces.
```

Some tap devices have assigned vlan tags which are still in use.

```
Jun 22 09:05:10 openqaworker6 ovs-vsctl[14369]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap6 tag=39
Jun 22 09:05:11 openqaworker6 ovs-vsctl[14373]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap1 tag=54
Jun 22 09:07:16 openqaworker6 ovs-vsctl[15211]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap4 tag=42
```

```
Jun 22 09:07:16 openqaworker6 ovs-vsctl[15214]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap68 tag=81
Jun 22 09:07:16 openqaworker6 ovs-vsctl[15217]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap132 tag=4
1
Jun 22 09:19:39 openqaworker6 ovs-vsctl[21624]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap0 tag=17
Jun 22 09:22:02 openqaworker6 ovs-vsctl[23163]: ovs|00001|vsctl|INFO|Called as ovs-vsctl remove port tap0 tag
17
Jun 22 09:24:39 openqaworker6 ovs-vsctl[24487]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap0 tag=39
Jun 22 09:24:40 openqaworker6 ovs-vsctl[24517]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap1 tag=41
Jun 22 09:24:41 openqaworker6 ovs-vsctl[24555]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap16 tag=17
Jun 22 09:24:41 openqaworker6 ovs-vsctl[24561]: ovs|00001|vsctl|INFO|Called as ovs-vsctl set port tap2 tag=82
```

#3 - 2018-08-07 07:27 - EDiGiacinto

- Status changed from New to Feedback
- Assignee set to EDiGiacinto
- Target version set to Current Sprint

We changed in <https://github.com/os-autoinst/os-autoinst/pull/1006> 'when' the vlan are untagged - now it's tied to the qemu process lifespan, but i need feedback on it (once it is deployed) since it will likely require other changes to cover edge cases.

#4 - 2018-08-29 06:19 - szarate

I think we haven't seen this one for a while

#5 - 2018-09-06 13:46 - thehejik

we still encounter with this, see <https://openqa.suse.de/tests/2029994/file/serial0.txt>

```
Sep 06 08:41:56 susetest wickedd[987]: eth0: IPv4 duplicate address 10.0.2.1 detected (in use by 52:54:00:12:03:c7)!
```

```
Sep 06 08:41:56 susetest wickedd[987]: __ni_rtnl_send_newroute(ipv4 0.0.0.0/0 via 10.0.2.2 dev eth0 type unicast table main scope universe
protocol boot): ni_nl_talk failed [Unspecific failure]
```

#6 - 2018-09-06 13:48 - asmorodskyi

another job with same failure - <https://openqa.suse.de/tests/2030015>

#7 - 2018-09-12 09:28 - coolo

- Status changed from Feedback to Resolved

(meeting protocol :)

We fixed 2 things:

- fixed the vlan assignment on scheduler side to make sure they don't collide between multimachine clusters
- fixed os-autoinst to kill the vlan when the qemu process is done

We think this will reduce the impact of the problem drastically, but the general solution to this race might still be needed. But we need to see how big the problem remains to be before we judge when to put worktime into it.

#8 - 2018-09-14 12:16 - coolo

- Target version changed from Current Sprint to Done