

## openQA Infrastructure - action #110467

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

coordination # 109659 (Blocked): [epic] More remote workers

### Establish reliable tap setup on ow14

2022-04-29 11:32 - mkittler

<b>Status:</b>	New	<b>Start date:</b>	2022-04-29
<b>Priority:</b>	Low	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	future		
<b>Description</b>			
After noticing connection problems with these two workers located in Prague the tap worker class has been removed again via salt.			
<b>Acceptance criteria</b>			
<ul style="list-style-type: none"><li><b>AC1:</b> ow14 can run multi-machine tests without running into connection issues (at least not more than usual)</li><li><b>AC2:</b> The multi-machine setup on ow14 recovers automatically after a temporary connection loss</li></ul>			
<b>Suggestions</b>			
<ul style="list-style-type: none"><li>The tap setup has been reverted by salt as the "tap" worker class has been removed. If one wanted to do further tap testing, the easiest solution is to use local salt/pillar checkouts under /srv and apply them via salt-call --local state.apply after removing the worker from the master via e.g. salt-key -d openqaworker14.qa.suse.cz. As far as I observed when working on <a href="#">#104970</a>, salt will configure everything needed for tap jobs to work and the configuration persisted after a reboot.</li><li>To spawn a parallel job cluster where some jobs are running on Nürnberg workers and some on Prague workers one can modify openqa-clone-job to print the parameters instead of posting them by putting print STDERR join(' ', map { "\$_=\$composed_params{\$_}" } sort keys %composed_params); exit 0; in post_jobs within CloneJob.pm. One can then modify the worker classes as needed and post the jobs manually.</li><li>One scenario affected by the so far unreliable tap setup is <a href="https://openqa.suse.de/tests/latest?arch=x86_64&amp;distri=sle&amp;flavor=Server-DVD-HPC-Incidents&amp;machine=64bit&amp;test=hpc_ALP_HA_mpih_mpi_supportserver&amp;version=15-SP1">https://openqa.suse.de/tests/latest?arch=x86_64&amp;distri=sle&amp;flavor=Server-DVD-HPC-Incidents&amp;machine=64bit&amp;test=hpc_ALP_HA_mpih_mpi_supportserver&amp;version=15-SP1</a>. So it can be cloned as mentioned in the previous point. (The scenario generally works when some jobs run on a mix of Nürnberg and Prague workers. I've tested that with ow3, ow10 and ow15. Only if the Prague workers are in some broken state the issue can be reproduced.)</li><li>I suppose one needs to research how gre tunnels generally behave on connection problems and how one can tweak the behavior.</li><li>Related threads on slack: <a href="https://suse.slack.com/archives/C02D16TCP99/p1651131920442009">https://suse.slack.com/archives/C02D16TCP99/p1651131920442009</a> (and before that there was <a href="https://suse.slack.com/archives/C02CANHLANP/p1651052375642909">https://suse.slack.com/archives/C02CANHLANP/p1651052375642909</a>)<ul style="list-style-type: none"><li>Note that the threads also stray in other issues like broken asset downloads on arm workers which are completely distinct issues.</li><li>We also have seen connection issues on other workers than ow14 and ow15 but way less.</li></ul></li><li>Try to create distinct "bubbles" of "tap"-enabled workers which are <i>not</i> interconnected. So e.g. the workers in Nürnberg and Prague would be in different "bubbles" and would therefore not be expected to be able to reach each other via gre tunnels as they use a distinct gre setup.<ul style="list-style-type: none"><li>The salt states should be able to establish distinct gre setups based on <i>some</i> configuration.</li><li>The openQA scheduler needed to be aware that a certain set of jobs must only run within a certain "bubble". We could use a distinct WORKER_CLASS but that setup would be rather static (we don't care in which specific bubble a parallel cluster runs, just that it is only run within the <i>one</i> bubble). openQA has already the concept of "vlan"s which jobs can be grouped by. <i>Maybe</i> that is helpful in this regard.</li></ul></li></ul>			
<b>Related issues:</b>			
Related to openQA Infrastructure - action #104970: Add two OSD workers (openq...		<b>Resolved</b>	<b>2022-01-17</b>

### History

#1 - 2022-04-29 11:32 - mkittler

- Tracker changed from coordination to action

#2 - 2022-04-29 11:35 - mkittler

- Project changed from openQA Project to openQA Infrastructure
- Category deleted (Feature requests)

**#3 - 2022-04-29 11:35 - mkittler**

- Related to action #104970: Add two OSD workers (openqaworker14+openqaworker15) specifically for sap-application testing size:M added

**#4 - 2022-04-29 12:41 - okurz**

- Target version set to Ready

**#5 - 2022-05-02 10:06 - okurz**

- Copied to action #110515: Command export feature for openqa-clone-job size:M added

**#6 - 2022-05-02 10:12 - okurz**

- Status changed from New to Blocked
- Assignee set to okurz

[#110515](#)

**#7 - 2022-05-03 14:06 - mkittler**

- Copied to deleted (action #110515: Command export feature for openqa-clone-job size:M)

**#8 - 2022-05-03 14:06 - mkittler**

- Status changed from Blocked to New

The ticket is no longer blocked by [#110515](#).

**#9 - 2022-05-05 08:26 - okurz**

- Subject changed from Establish reliable tap setup on ow14+15 to Establish reliable tap setup on ow14
- Assignee deleted (okurz)

**#10 - 2022-05-05 08:27 - okurz**

- Description updated

**#11 - 2022-05-06 10:19 - mkittler**

- Description updated

Added suggestion about distinct "bubbles" of "tap"-enabled workers.

**#12 - 2022-05-06 10:26 - mkittler**

- Description updated

**#13 - 2022-05-16 15:02 - okurz**

- Priority changed from Normal to Low
- Target version changed from Ready to future