

openQA Project - action #98841

qemu randomly fails to start on QA-Power8-5-kvm auto_review:"(?s)QA-Power8.*Failed to allocate KVM HPT of order 25 (try smaller maxmem?): Cannot allocate memory":retry size:M

2021-09-17 15:49 - MDoucha

Status:	Resolved	Start date:	2021-09-17
Priority:	Normal	Due date:	2021-10-28
Assignee:	Xiaojing_liu	% Done:	0%
Category:	Feature requests	Estimated time:	30.00 hours
Target version:	Ready		
Difficulty:			

Description

Observation

A few LTP jobs have failed to start today due to qemu error on QA-Power8-5-kvm worker:

<https://openqa.suse.de/tests/7138972>

<https://openqa.suse.de/tests/7149857>

<https://openqa.suse.de/tests/7153989>

All of them have the following output in autoinst-log.txt:

```
[2021-09-17T16:24:29.803 CEST] [info] ::: backend::baseclass::die_handler: Backend process died, backend errors are reported below in the following lines:
```

```
QEMU terminated before QMP connection could be established. Check for errors below
```

```
[2021-09-17T16:24:29.804 CEST] [info] ::: OpenQA::Qemu::Proc::save_state: Saving QEMU state to qemu_state.json
```

```
[2021-09-17T16:24:29.805 CEST] [debug] Passing remaining frames to the video encoder
```

```
[2021-09-17T16:24:29.805 CEST] [debug] Waiting for video encoder to finalize the video
```

```
[2021-09-17T16:24:29.805 CEST] [debug] The built-in video encoder (pid 110385) terminated
```

```
[2021-09-17T16:24:29.807 CEST] [debug] QEMU: QEMU emulator version 4.2.1 (openSUSE Leap 15.2)
```

```
[2021-09-17T16:24:29.807 CEST] [debug] QEMU: Copyright (c) 2003-2019 Fabrice Bellard and the QEMU Project developers
```

```
[2021-09-17T16:24:29.807 CEST] [warn] !!! : qemu-system-ppc64: Failed to allocate KVM HPT of order 25 (try smaller maxmem?): Cannot allocate memory
```

Problem

QA-Power8-5-kvm has 256GB RAM.

<https://monitor.qa.suse.de/d/WDQA-Power8-5-kvm/worker-dashboard-qa-power8-5-kvm?viewPanel=12054&orgId=1&from=1631765162464&to=1632085860553> shows that some memory was used during the period when the test failed but nothing that should explain the inability to allocate the memory for the qemu VM. In the system journal there is

```
Sep 17 16:24:28 QA-Power8-5-kvm worker[88148]: [debug] [pid:88148] REST-API call: POST http://openqa.suse.de/api/v1/jobs/7125263/status
```

```
Sep 17 16:24:29 QA-Power8-5-kvm worker[104911]: [info] [pid:108741] sle-15-SP4-ppc64le-Build36.1-HA-BV.qcow2: Processing chunk 501/5812, avg. speed ~976.562 KiB/s
```

```
Sep 17 16:24:29 QA-Power8-5-kvm worker[101413]: [debug] [pid:102598] Uploading artefact mq_timedreceiving_15-1-2.txt
```

```
Sep 17 16:24:29 QA-Power8-5-kvm worker[96737]: [debug] [pid:96737] REST-API call: POST http://openqa.suse.de/api/v1/jobs/7125265/status
```

```
Sep 17 16:24:29 QA-Power8-5-kvm worker[109458]: [debug] [pid:110336] Uploading artefact bootloader_start-15.txt
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: 23 callbacks suppressed
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf00, 3cb100) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf04, 3cb104) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf08, 3cb108) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf0c, 3cb10c) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf10, 3cb110) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf14, 3cb114) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf18, 3cb118) PFNs busy
```

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf1c, 3cb11c) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf20, 3cb120) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf24, 3cb124) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: cma: cma_alloc: alloc failed, req-size: 512 pages, ret: -1
6
Sep 17 16:24:30 QA-Power8-5-kvm worker[88148]: [debug] [pid:88148] Upload concluded (at wait_child
ren)
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [info] [pid:109557] Isotovideo exit status: 1
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [debug] [pid:109557] Stopping job 7153989 from ope
nqa.suse.de: 07153989-sle-15-SP3-Server-DVD-Incidents-Kernel-KOTD-ppc64le-Build5.3.18-302.1.g31699
3b-ltp_crashme@ppc64le-virtio - reason: died
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [debug] [pid:109557] REST-API call: POST http://op
enqa.suse.de/api/v1/jobs/7153989/status
Sep 17 16:24:30 QA-Power8-5-kvm worker[101413]: [debug] [pid:102598] Uploading artefact mq_timedre
ceive_7-1-2.txt
```

in particular the messages

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf20, 3cb120) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf24, 3cb124) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: cma: cma_alloc: alloc failed, req-size: 512 pages, ret: -1
6
```

so an allocation failure. We could report a bug about this but because KVM on SUSE with Power8 is unsupported so I don't expect any success.

We likely need to accept such issues and trigger a restart automatically by openQA.

Acceptance criteria

- **AC1:** qemu ppc64le allocation errors cause automatic job retriggers by openQA

Suggestions

- Catch the error and make it "Incomplete"
- Restart the incomplete job
- Make openQA automatically detect the issue and trigger restart, e.g. based on <https://github.com/os-autoinst/openQA/blob/master/etc/openqa/openqa.ini#L76>

History

#1 - 2021-09-20 09:35 - okurz

- Target version set to Ready

#2 - 2021-09-20 19:11 - okurz

- Project changed from openQA Infrastructure to openQA Project

- Subject changed from qemu randomly fails to start on QA-Power8-5-kvm to qemu randomly fails to start on QA-Power8-5-kvm auto_review:"(?s)QA-Power8.*Failed to allocate KVM HPT of order 25 (try smaller maxmem?): Cannot allocate memory":retry

- Description updated

- Category set to Feature requests

QA-Power8-5-kvm has 256GB RAM.

<https://monitor.qa.suse.de/d/WDOQA-Power8-5-kvm/worker-dashboard-qa-power8-5-kvm?viewPanel=12054&orgId=1&from=1631765162464&to=1632085860553> shows that some memory was used during the period when the test failed but nothing that should explain the inability to allocate the memory for the qemu VM. In the system journal there is

```
Sep 17 16:24:28 QA-Power8-5-kvm worker[88148]: [debug] [pid:88148] REST-API call: POST http://openqa.suse.de/a
pi/v1/jobs/7125263/status
Sep 17 16:24:29 QA-Power8-5-kvm worker[104911]: [info] [pid:108741] sle-15-SP4-ppc64le-Build36.1-HA-BV.qcow2:
Processing chunk 501/5812, avg. speed ~976.562 KiB/s
Sep 17 16:24:29 QA-Power8-5-kvm worker[101413]: [debug] [pid:102598] Uploading artefact mq_timedreceive_15-1-2
.txt
Sep 17 16:24:29 QA-Power8-5-kvm worker[96737]: [debug] [pid:96737] REST-API call: POST http://openqa.suse.de/a
pi/v1/jobs/7125265/status
Sep 17 16:24:29 QA-Power8-5-kvm worker[109458]: [debug] [pid:110336] Uploading artefact bootloader_start-15.tx
```

t
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: 23 callbacks suppressed
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf00, 3cb100) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf04, 3cb104) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf08, 3cb108) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf0c, 3cb10c) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf10, 3cb110) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf14, 3cb114) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf18, 3cb118) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf1c, 3cb11c) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf20, 3cb120) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf24, 3cb124) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: cma: cma_alloc: alloc failed, req-size: 512 pages, ret: -16
Sep 17 16:24:30 QA-Power8-5-kvm worker[88148]: [debug] [pid:88148] Upload concluded (at wait_children)
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [info] [pid:109557] Isotovideo exit status: 1
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [debug] [pid:109557] Stopping job 7153989 from openqa.suse.de: 07153989-sle-15-SP3-Server-DVD-Incidents-Kernel-KOTD-ppc64le-Build5.3.18-302.1.g316993b-ltp_crashme@ppc64le-virtio - reason: died
Sep 17 16:24:30 QA-Power8-5-kvm worker[109557]: [debug] [pid:109557] REST-API call: POST http://openqa.suse.de/api/v1/jobs/7153989/status
Sep 17 16:24:30 QA-Power8-5-kvm worker[101413]: [debug] [pid:102598] Uploading artefact mq_timedreceive_7-1-2.txt

in particular the messages

```
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf20, 3cb120) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: alloc_contig_range: [3caf24, 3cb124) PFNs busy
Sep 17 16:24:29 QA-Power8-5-kvm kernel: cma: cma_alloc: alloc failed, req-size: 512 pages, ret: -16
```

so an allocation failure. We could report a bug about this but because KVM on SUSE with Power8 is unsupported so I don't expect any success.

We likely need to accept such issues and trigger a restart automatically by openQA.

#3 - 2021-10-07 09:46 - cdywan

- Subject changed from *qemu randomly fails to start on QA-Power8-5-kvm auto_review:"(?s)QA-Power8.*Failed to allocate KVM HPT of order 25 (try smaller maxmem?): Cannot allocate memory":retry to qemu randomly fails to start on QA-Power8-5-kvm auto_review:"(?s)QA-Power8.*Failed to allocate KVM HPT of order 25 (try smaller maxmem?): Cannot allocate memory":retry size:M*

- Description updated

- Status changed from *New* to *Workable*

#4 - 2021-10-11 01:50 - Xiaoqing_liu

- Assignee set to *Xiaoqing_liu*

#5 - 2021-10-11 08:44 - cdywan

Discussed briefly in the daily. The job is already *Incomplete*, so we only need to handle the restart

#6 - 2021-10-13 09:11 - Xiaoqing_liu

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

PR: <https://github.com/os-autoinst/openQA/pull/4298>

#7 - 2021-10-14 04:10 - openqa_review

- Due date set to *2021-10-28*

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

#8 - 2021-10-21 03:02 - Xiaoqing_liu

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

The PR got merged.

#9 - 2021-10-22 11:31 - Xiaoqing_liu

- Status changed from *Feedback* to *Resolved*

We could track this ticket by seeing monitor results.

#10 - 2021-10-27 01:28 - Xiaojing_liu

- Estimated time set to 30.00 h