

openQA Project - action #94850

QEMU 6.0 fails to start if job has QEMU_NUMA=1

2021-06-29 10:05 - MDoucha

Status:	Resolved	Start date:	2021-06-29
Priority:	Normal	Due date:	
Assignee:	okurz	% Done:	0%
Category:	Feature requests	Estimated time:	0.00 hour
Target version:	Ready		
Difficulty:			

Description

Observation

Originally reported by pvorel relating to <http://quasar.suse.cz/tests/6860#>

It appears that QEMU 6.0 requires explicit memory assignment to NUMA nodes using `-numa node,mem=...` command line arguments. OpenQA jobs with `QEMU_NUMA=1` currently fail to start on Tumbleweed because `os-autoinst` omits the `mem` option, which leads to the following error:

```
[2021-06-28T13:26:00.308 CEST] [debug] starting: /usr/bin/qemu-system-x86_64 -only-migratable -chardev ringbuf,id=serial0,logfile=serial0,logappend=on -serial chardev:serial0 -audiodev none,id=snd0 -device intel-hda -device hda-output,audiodev=snd0 -global isa-fdc.fdtypeA=none -m 2048 -cpu qemu64 -netdev user,id=qanet0 -device virtio-net,netdev=qanet0,mac=52:54:00:12:34:56 -boot order=c -device usb-ehci -device usb-tablet -smp 2 -numa node,nodeid=0 -numa node,nodeid=1 -enable-kvm -no-shutdown -vnc :91,share=force-shared -device virtio-serial -chardev pipe,id=virtio_console,path=virtio_console,logfile=virtio_console.log,logappend=on -device virtconsole,chardev=virtio_console,name=org.openqa.console.virtio_console -chardev socket,path=qmp_socket,server,nowait,id=qmp_socket,logfile=qmp_socket.log,logappend=on -qmp chardev:qmp_socket -S -device virtio-scsi-pci,id=scsi0 -blockdev driver=file,node-name=hd0-overlay0-file,filename=/var/lib/openqa/pool/1/raid/hd0-overlay0,cache.no-flush=on -blockdev driver=qcow2,node-name=hd0-overlay0,file=hd0-overlay0-file,cache.no-flush=on -device virtio-blk,id=hd0-device,drive=hd0-overlay0,bootindex=0,serial=hd0
[2021-06-28T13:26:00.312 CEST] [debug] Waiting for 0 attempts
[2021-06-28T13:26:00.423 CEST] [debug] Waiting for 1 attempts
[2021-06-28T13:26:00.424 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 at /usr/lib/os-autoinst/OpenQA/Qemu/Proc.pm line 453.

[2021-06-28T13:26:00.424 CEST] [info] ::: OpenQA::Qemu::Proc::save_state: Saving QEMU state to qemu_state.json
[2021-06-28T13:26:00.424 CEST] [debug] Passing remaining frames to the video encoder
[2021-06-28T13:26:00.426 CEST] [debug] Waiting for video encoder to finalize the video
[2021-06-28T13:26:00.426 CEST] [debug] The built-in video encoder (pid 22047) terminated
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: QEMU emulator version 6.0.0 (openSUSE Tumbleweed)
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: Copyright (c) 2003-2021 Fabrice Bellard and the QEMU Project developers
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: qemu-system-x86_64: -chardev socket,path=qmp_socket,server,nowait,id=qmp_socket,logfile=qmp_socket.log,logappend=on: warning: short-form boolean option 'server' deprecated
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: Please use server=on instead
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: qemu-system-x86_64: -chardev socket,path=qmp_socket,server,nowait,id=qmp_socket,logfile=qmp_socket.log,logappend=on: warning: short-form boolean option 'nowait' deprecated
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: Please use wait=off instead
[2021-06-28T13:26:00.427 CEST] [debug] QEMU: qemu-system-x86_64: total memory for NUMA nodes (0x0) should equal RAM size (0x80000000)
```

Acceptance criteria

- **AC1:** `os-autoinst` continues to start for `qemu >= 6.0` and `QEMU_NUMA=1`
- **AC2:** `os-autoinst` jobs still work fine on `qemu < 6.0` regardless of `QEMU_NUMA` setting

History

#1 - 2021-06-29 10:18 - okurz

- *Target version set to Ready*

#2 - 2021-06-29 10:22 - okurz

- *Description updated*

- *Category changed from Concrete Bugs to Feature requests*

- *Status changed from New to In Progress*

- *Assignee set to okurz*

#3 - 2021-06-29 10:50 - okurz

- *Status changed from In Progress to Feedback*

<https://github.com/os-autoinst/os-autoinst/pull/1707>

#4 - 2021-07-01 16:03 - MDoucha

I've tested the new os-autoinst package on my Tumbleweed machine, jobs with QEMU_NUMA=1 are booting again on QEMU 6.0.

#5 - 2021-07-01 18:44 - okurz

- *Status changed from Feedback to Resolved*

Thanks for verification. As in production we run openSUSE Leap 15.2 we are not affected by this so with the personal verification I consider this enough to resolve the ticket