

openQA Infrastructure - action #20914

[tools] configure vm settings for workers with rotating discs

2017-07-28 12:59 - coolo

Status:	Resolved	Start date:	2017-07-28
Priority:	Normal	Due date:	2019-11-05
Assignee:	okurz	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Current Sprint		
Description			
Especially aarch64 machines are too slow syncing qemu, so we need to tweak their configs in salt			
This will cost performance - and possibly making the 'HMP timeout' issue more prominent, but it will also make the needling matching more predictable.			
Jan Kara's recommendation is: dirty_bytes to 200000000 (~200 MB) and dirty_background_bytes to 50000000 (~50 MB).			
after the experiments in https://github.com/os-autoinst/os-autoinst/pull/664			
We only need this for the HDD hosts, having it on NVME shouldn't hurt - but I can't really say			
Related issues:			
Related to openQA Infrastructure - action #58805: [infra]Severe storage perfo...		Resolved	2019-10-29
Related to openQA Tests - action #50615: [functional][y] test fails in await_...		Resolved	2019-04-22

History

#1 - 2018-10-16 09:18 - coolo

- Project changed from openQA Tests to openQA Infrastructure
- Category deleted (Infrastructure)

#2 - 2018-11-20 13:54 - nicksinger

- Status changed from New to Workable

#3 - 2019-09-24 19:16 - okurz

[coolo](#) do you think we should still try to tinker with these variables? I don't think the mentioned problems are relevant anymore but of course we can still try improve based on vm options.

#4 - 2019-09-25 07:55 - coolo

What makes you think linux's memory management got any better since?

#5 - 2019-10-15 12:47 - okurz

- Status changed from Workable to Feedback
- Assignee set to okurz
- Target version set to Current Sprint

https://gitlab.suse.de/openqa/salt-states-openqa/merge_requests/203

#6 - 2019-10-22 13:42 - okurz

- Due date set to 2019-11-05

merged. Let's monitor if it has any measurable impact.

#7 - 2019-10-30 12:06 - coolo

- Related to action #58805: [infra]Severe storage performance issue on openqa.suse.de workers added

#8 - 2019-10-30 12:07 - coolo

It had - I'm going to increase it to 10%/5% again. This is still 50% of the default, but way above the current settings

#9 - 2019-10-30 12:08 - coolo

I asked the DLs for 5 SSDs, let's see :)

#10 - 2019-10-30 20:45 - okurz

You did https://gitlab.suse.de/openqa/salt-states-openqa/merge_requests/215 and called it "Increase the dirty buffer size" whoever I believe you are actually *decreasing* it as the values are lower than default.

I have good experience with the following:

```
# https://askubuntu.com/questions/157793/why-is-swap-being-used-even-though-i-have-plenty-of-free-ram
# https://askubuntu.com/questions/440326/how-can-i-turn-off-swap-permanently
# https://superuser.com/questions/1115983/prevent-system-freeze-unresponsiveness-due-to-swapping-run-away-memory-usage
vm.dirty_background_ratio = 5
vm.dirty_ratio = 80
# okurz: 2019-01-04: Trying to prevent even more stuttering
# vm.swappiness = 10
# https://rudd-o.com/linux-and-free-software/tales-from-responsivenessland-why-linux-feels-slow-and-how-to-fix-that
vm.swappiness = 1
# did not actually experiment with finding a good value, just took the one from the above webpage
vm.vfs_cache_pressure = 50
```

As an alternative we can say whenever we hit problems due to this we need to simply buy more RAM.

WDYT?

#11 - 2019-10-31 05:17 - okurz

- Related to action #50615: `[functional][y]` test fails in `await_install` - does not catch `rebootnow` added

#12 - 2019-11-04 05:53 - coolo

you don't understand the problem I'm afraid. this has nothing to do with RAM nor with swap.

#13 - 2019-11-04 15:46 - okurz

ok maybe I was misleading with mentioning the part about swap or thrashing. It's not about memory depletion for sure. So let me simply ask: Did you not *decrease* the values below default now?

#14 - 2019-11-04 19:15 - coolo

the default is 10% of memory which is about 26GB - our initial hit was at 200MB (which is less than 1% of default), which was too small. Now we're at 5% of memory, which is somewhere in the middle

#15 - 2019-11-04 21:40 - okurz

- Status changed from *Feedback* to *Resolved*

Exactly. Anyway, I guess we can call this solved then. Adjusting the values is easy now and we can also make it smart, when necessary