Subject: How to enable swap for CT when swap on OpenVZ 7 host machine is disabled? Posted by iblinder on Wed, 24 Oct 2018 08:54:32 GMT View Forum Message <> Reply to Message

Hello.

We can disable swap on OpenVZ 6 host machine, but swap in containers works.

When we turn off swap on VZ 7 host machine, swap in containers not works. Why we can't to disable swap on host if swap for containers is in host machine memory?

Thanks.

Subject: Re: How to enable swap for CT when swap on OpenVZ 7 host machine is disabled?

Posted by khorenko on Mon, 08 Jun 2020 09:41:45 GMT View Forum Message <> Reply to Message

Yes, i know the thread is old.

i'm writing here because this page is mentioned in the OpenVZ bug https://bugs.openvz.org/browse/OVZ-7072

Quote from the bug:

Quote:

Guys, you've described vSwap behavior in vz7 correctly:

* if some page goes into virtual swap inside a Container, a new page is allocated in physical swap on the Hardware Node. Always.

* if there is free RAM on the host, the page is NOT written to physical swap on the host, but saved in free host RAM (tswap area)

* tswap - is a cache, so on a memory pressure (on host) it must free its pages, so it MUST have some space to write the data - that's why swap entry is allocated in advance - it guarantees all tswap entries can be freed any time later

* important note: once again - physical swap entries on host are allocated if a CT uses vswap, but real data is NOT written to the physical disk in case there is free RAM on host

On vz6 we had similar behavior initially but later had to rework so vswap does not allocate physical swap entries.

This was done because online migration required it: technically in vz6 online migration marks all memory pages as "in swap", thus is required lot of physical swap - not only for CT vswap usage, but for all memory usage as well. And sometimes online migration fails to to luck of physical swap.

In vz7 CRIU online migration work in another way, so no need for such hack anymore and thus we do not plan to change this behavior.

Disks are big now, so increasing swap space on the host should not be a big problem (once again - there is no real swapin/swapout if there is free RAM on host).

Or just don't set up vSwap for Containers: if host has free RAM, the performance is the same (if CT has RAM=4Gb vs CT has RAM 2Gb + vSwap 2Gb),

and in case host does not have free RAM (so vSwap does not work at all, host physical swap is used), host memory management knows better what memory to put into physical swap.

Page 2 of 2 ---- Generated from OpenVZ Forum