
Subject: Re: [RFC][PATCH] another swap controller for cgroup

Posted by [yamamoto](#) on Tue, 08 Apr 2008 03:29:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

> YAMAMOTO Takashi wrote:

> > hi,

> >

> > i tried to reproduce the large swap cache issue, but no luck.

> > can you provide a little more detailed instruction?

> >

> This issue also happens on generic 2.6.25-rc3-mm1

> (with limiting only memory), so I think this issue is not

> related to your patch.

> I'm investigating this issue too.

>

> Below is my environment and how to reproduce.

>

> - System

> full virtualized xen guest based on RHEL5.1(x86_64).

> CPU: 2

> memory: 2GB

> swap: 1GB

> A config of the running kernel(2.6.25-rc3-mm1 with your patch)

> is attached.

>

> - how to reproduce

> - change swappiness to 100

>

> echo 100 >/proc/sys/vm/swappiness

>

> - mount cgroup fs

>

> # mount -t cgroup -o memory,swap none /cgroup

>

> - make cgroup for test

>

> # mkdir /cgroup/02

> # echo -n 64M >/cgroup/02/memory.limit_in_bytes

> # echo -n `expr 128 * 1024 * 1024` >/cgroup/02/swap.limit_in_bytes

>

> - run test

>

> # echo \$\$ >/cgroup/02/tasks

> # while true; do make clean; make -j2; done

>

> In other terminals, I run some monitoring processes, top,

> "tail -f /var/log/messages", and displaying *.usage_in_bytes

> every seconds.

>
>
> Thanks,
> Daisuke Nishimura.

what i tried was essentially same.
for me, once `vm_swap_full()` got true, swap cache stopped growing as expected.

<http://people.valinux.co.jp/~yamamoto/swap.png>

it was taken by running
while :;do swapon -s|tail -1;sleep 1;done > foo
in an unlimited cgroup, and then plotted by gnuplot.
plot "foo" u 4

as my system has 1GB swap configured, the `vm_swap_full()` threshold is
around 500MB.

YAMAMOTO Takashi

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
