
Subject: Re: [RFC][PATCH] another swap controller for cgroup
Posted by [Balbir Singh](#) on Mon, 24 Mar 2008 12:22:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daisuke Nishimura wrote:

> Daisuke Nishimura wrote:

>> Hi, Yamamoto-san.

>>

>> I'm reviewing and testing your patch now.

>>

>

> In building kernel infinitely(in a cgroup of

> memory.limit=64M and swap.limit=128M, with swappiness=100),

> almost all of the swap (1GB) is consumed as swap cache

> after a day or so.

> As a result, processes are occasionally OOM-killed even when

> the swap.usage of the group doesn't exceed the limit.

>

> I don't know why the swap cache uses up swap space.

> I will test whether a similar issue happens without your patch.

> Do you have any thoughts?

>

>

> BTW, I think that it would be better, in the sence of

> isolating memory resource, if there is a framework

> to limit the usage of swap cache.

We had this earlier, but dropped it later due to issues related to swap
readahead and assigning the pages to the correct cgroup.

--

Warm Regards,

Balbir Singh

Linux Technology Center

IBM, ISTL

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
