Subject: Re: [PATCH 0/4] swapcgroup(v2)
Posted by Rik van Riel on Fri, 23 May 2008 02:26:55 GMT
View Forum Message <> Reply to Message

On Thu, 22 May 2008 15:13:41 +0900
Daisuke Nishimura <nishimura@mxp.nes.nec.co.jp> wrote:

> I updated my swapcgroup patch.

I do not understand why this is useful.

With the other cgroup resource controllers, once a process group reaches its limit, it is limited or punished in some way. For example, when it goes over its RSS limit, memory is taken away.

However, once a cgroup reaches its swap limit, it is rewarded, by allowing more of its pages to stay resident in RAM, instead of having them swapped out.

This, in turn, will cause the VM to evict pages from other, better behaving groups. In short, the cgroup that has "misbehaved" by reaching its limit causes other cgroups to get punished.

Even worse is that a cgroup has NO CONTROL over how much of its memory is kept in RAM and how much is swapped out. This kind of decision is made on a system-wide basis by the kernel, dependent on what other processes in the system are doing. There also is no easy way for a cgroup to reduce its swap use, unlike with other resources.

In what scenario would you use a resource controller that rewards a group for reaching its limit?

How can the cgroup swap space controller help sysadmins achieve performance or fairness goals on a system?

--
All rights reversed.

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