
Subject: Re: [RFC][PATCH] memory cgroup enhancements updated [7/10]
RSS/CACHE failcnt

Posted by [KAMEZAWA Hiroyuki](#) on Fri, 19 Oct 2007 22:33:38 GMT

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

On Fri, 19 Oct 2007 09:54:23 -0700

"Paul Menage" <menage@google.com> wrote:

> On 10/19/07, KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:

> > cgroup's resource has "failure" counter. But I think memory cgroup

> > has 2 types of failure

> > - failure of cache

> > - failure of RSS

>

> Why do you think these are significantly different? Also, what kind of

> failure do we get when a task tries to fault on a mapping for a

> file-backed VMA? That's both RSS (since it's due to be mapped into the

> process' memory) and pagecache (assuming a cache-backed filesystem).

>

One is for checking how memory cgroup works.

When memory reaches limit and amount of CACHE doesn't change but

CACHE.failure increases, maybe CACHE is rotating.

One is for debug, I'd like to check swapiness <-> RSS.failure:CACHE.failure
relationship. It's ok to turn these params to be DEBUG option.

File-backed VMA gets RSS.failure.

Thanks,

-Kame

Containers mailing list

Containers@lists.linux-foundation.org

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