
Subject: Re: [PATCH 3/5] page_cgroup: make page tracking available for blkio
Posted by [Andrea Righi](#) on Tue, 22 Feb 2011 23:48:33 GMT

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

On Tue, Feb 22, 2011 at 04:27:29PM -0700, Jonathan Corbet wrote:

> On Wed, 23 Feb 2011 00:01:47 +0100

> Andrea Righi <arighi@develer.com> wrote:

>

> > > My immediate observation is that you're not really tracking the "owner"

> > > here - you're tracking an opaque 16-bit token known only to the block

> > > controller in a field which - if changed by anybody other than the block

> > > controller - will lead to mayhem in the block controller. I think it

> > > might be clearer - and safer - to say "blkcg" or some such instead of

> > > "owner" here.

> >

> > Basically the idea here was to be as generic as possible and make this

> > feature potentially available also to other subsystems, so that cgroup

> > subsystems may represent whatever they want with the 16-bit token.

> > However, no more than a single subsystem may be able to use this feature

> > at the same time.

>

> That makes me nervous; it can't really be used that way unless we want to

> say that certain controllers are fundamentally incompatible and can't be

> allowed to play together. For whatever my \$0.02 are worth (given the

> state of the US dollar, that's not a whole lot), I'd suggest keeping the

> current mechanism, but make it clear that it belongs to your controller.

> If and when another controller comes along with a need for similar

> functionality, somebody can worry about making it more general.

OK, I understand. I'll use "blkio" instead of "owner". Also because I wouldn't like to introduce additional logic and overhead to check if two controllers are using this feature at the same time. Better to hard-code this information in the name of the functions.

Probably the most generic solution is the one that you suggested: replace the mem_cgroup with a pointer to css_set. I'll also try to investigate this way.

Thanks,
-Andrea

Containers mailing list

Containers@lists.linux-foundation.org

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