
Subject: Re: [PATCH 0/8] RSS controller based on process containers (v2)
Posted by [Vaidyanathan Srinivas](#) on Thu, 19 Apr 2007 05:37:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Pavel Emelianov wrote:

> Peter Zijlstra wrote:

>> *ugh* /me no like.

>>

>> The basic premises seems to be that we can track page owners perfectly

>> (although this patch set does not yet do so), through get/release

>

> It looks like you have examined the patches not very carefully

> before concluding this. These patches DO track page owners.

>

> I know that a page may be shared among several containers and

> thus have many owners so we should track all of them. This is

> exactly what we decided not to do half-a-year ago.

>

> Page sharing accounting is performed in OpenVZ beancounters, and

> this functionality will be pushed to mainline after this simple

> container.

>

>> operations (on _mapcount).

>>

>> This is simply not true for unmapped pagecache pages. Those receive no

>> 'release' event; (the usage by find_get_page() could be seen as 'get').

>

> These patches concern the mapped pagecache only. Unmapped pagecache

> control is out of the scope of it since we do not want one container

> to track all the resources.

Unmapped pagecache control and swapcache control is part of independent pagecache controller that is being developed. Initial version was posted at <http://lkml.org/lkml/2007/3/06/51>

I plan to post a new version based on this patchset in a couple of days.

--Vaidy

>> Also, you don't seem to balance the active/inactive scanning on a per

>> container basis. This skews the per container working set logic.

>

> This is not true. Balbir sent a patch to the first version of this

> container that added active/inactive balancing to the container.

> I have included this (a bit reworked) patch into this version and

> pointed this fact in the zeroth letter.

>

[snip]
