

---

Subject: Re: [-mm PATCH 0/7] Memory controller introduction  
Posted by [Pavel Emelianov](#) on Thu, 05 Jul 2007 09:14:00 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Balbir Singh wrote:

- > Resending with the patch numbering fixed and linux-mm copied
- >
- > This patchset implements another version of the memory controller. These
- > patches have been through a big churn, the first set of patches were posted
- > last year and earlier this year at
- > <http://lkml.org/lkml/2007/2/19/10>
- >
- > Ever since, the RSS controller has been through four revisions, the latest
- > one being
- > <http://lwn.net/Articles/236817/>
- >
- > This patchset draws from the patches listed above and from some of the
- > contents of the patches posted by Vaidyanathan for page cache control.
- > <http://lkml.org/lkml/2007/6/20/92>
- >
- > Pavel, Vaidy could you look at the patches and add your signed off by
- > where relevant?

As far as I remember at OLS we decided to implement per-zone RLU lists and reuse the lru lock as well. This will remove all the problems with per-container lists inconsistency.

Separate limits for RSS and RSS+pagecache are also a must.

BTW, if you send smb. else's patches you may include a 'From: xxx' line into the letter to address the original author.

- > At OLS, the resource management BOF, it was discussed that we need to manage
- > RSS and unmapped page cache together. This patchset is a step towards that
- >
- > TODO's
- >
- > 1. Add memory controller water mark support. Reclaim on high water mark
- > 2. Add support for shrinking on limit change
- > 3. Add per zone per container LRU lists
- > 4. Make page\_referenced() container aware
- > 5. Figure out a better CLUI for the controller
- >
- > In case you have been using/testing the RSS controller, you'll find that
- > this controller works slower than the RSS controller. The reason being
- > that both swap cache and page cache is accounted for, so pages do go
- > out to swap upon reclaim (they cannot live in the swap cache).
- >

> I've test compiled the framework without the controller enabled, tested  
> the code on UML and minimally on a power box.  
>  
> Any test output, feedback, comments, suggestions are welcome!  
>  
> series  
>  
> res\_counters\_infra.patch  
> mem-control-setup.patch  
> mem-control-accounting-setup.patch  
> mem-control-accounting.patch  
> mem-control-task-migration.patch  
> mem-control-lru-and-reclaim.patch  
> mem-control-out-of-memory.patch  
>

---

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

---