
Subject: Re: [RFC] mm-controller

Posted by [Balbir Singh](#) on Mon, 25 Jun 2007 17:33:24 GMT

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

Peter Zijlstra wrote:

> On Fri, 2007-06-22 at 22:05 +0530, Vaidyanathan Srinivasan wrote:

>

>> Merging both limits will eliminate the issue, however we would need
>> individual limits for pagecache and RSS for better control. There are
>> use cases for pagecache_limit alone without RSS_limit like the case of
>> database application using direct IO, backup applications and
>> streaming applications that does not make good use of pagecache.

>

> I'm aware that some people want this. However we rejected adding a
> pagecache limit to the kernel proper on grounds that reclaim should do a
> better job.

>

> And now we're sneaking it in the backdoor.

>

We'll we are trying to provide a memory controller and page cache is a part of memory. The page reclaimer does treat page cache separately. Isn't this approach better than simply extending the vm_swappiness to per container?

> If we're going to do this, get it in the kernel proper first.

>

I'm open to this. There were several patches to do this. We can do this by splitting the LRU list to mapped and unmapped pages or by trying to balance the page cache by tracking it's usage.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
