

---

Subject: Re: [RFC][PATCH][3/4] Add reclaim support  
Posted by [Balbir Singh](#) on Mon, 19 Feb 2007 11:16:33 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Andrew Morton wrote:

> On Mon, 19 Feb 2007 16:20:53 +0530 Balbir Singh <balbir@in.ibm.com> wrote:

>

>>>> + \* so, is the container over it's limit. Returns 1 if the container is above

>>>> + \* its limit.

>>>> + \*/

>>>> +int memctlr\_mm\_overlimit(struct mm\_struct \*mm, void \*sc\_cont)

>>>> +{

>>>> + struct container \*cont;

>>>> + struct memctlr \*mem;

>>>> + long usage, limit;

>>>> + int ret = 1;

>>>> +

>>>> + if (!sc\_cont)

>>>> + goto out;

>>>> +

>>>> + read\_lock(&mm->container\_lock);

>>>> + cont = mm->container;

>>>> +

>>>> + /\*

>>>> + \* Regular reclaim, let it proceed as usual

>>>> + \*/

>>>> + if (!sc\_cont)

>>>> + goto out;

>>>> +

>>>> + ret = 0;

>>>> + if (cont != sc\_cont)

>>>> + goto out;

>>>> +

>>>> + mem = memctlr\_from\_cont(cont);

>>>> + usage = atomic\_long\_read(&mem->counter.usage);

>>>> + limit = atomic\_long\_read(&mem->counter.limit);

>>>> + if (limit && (usage > limit))

>>>> + ret = 1;

>>>> +out:

>>>> + read\_unlock(&mm->container\_lock);

>>>> + return ret;

>>>> +}

>>> hm, I wonder how much additional lock traffic all this adds.

>>>

>> It's a read\_lock() and most of the locks are read\_locks

>> which allow for concurrent access, until the container

>> changes or goes away

>

> read\_lock isn't free, and I suspect we're calling this function pretty  
> often (every pagefault?) It'll be measurable on some workloads, on some  
> hardware.  
>  
> It probably won't be terribly bad because each lock-taking is associated  
> with a clear\_page(). But still, if there's any possibility of lightening  
> the locking up, now is the time to think about it.  
>

Yes, good point. I'll revisit to see if barriers can replace the locking  
or if the locking is required at all?

```
>>>> @@ -66,6 +67,9 @@ struct scan_control {
>>>> int swappiness;
>>>>
>>>> int all_unreclaimable;
>>>> +
>>>> + void *container; /* Used by containers for reclaiming */
>>>> + /* pages when the limit is exceeded */
>>>> };
>>> eww. Why void*?
>>>
>> I did not want to expose struct container in mm/vmscan.c.
>
> It's already there, via rmap.h
>
```

Yes, true

```
>> An additional
>> thought was that no matter what container goes in the field would be
>> useful for reclaim.
>
> Am having trouble parsing that sentence ;)
>
>
```

The thought was that irrespective of the infrastructure that goes in  
having an entry for reclaim in scan\_control would be useful. I guess  
the name exposes what the type tries to hide :-)

--

Warm Regards,  
Balbir Singh

---