
Subject: [RFC][for -mm] memory controller enhancements for reclaiming take2 [8/8]
wake up waiters at unchage

Posted by [KAMEZAWA Hiroyuki](#) on Mon, 03 Dec 2007 09:45:36 GMT

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

Throttling direct reclaim reduces the sytem load. But waiters are only waken up if someone finish try_to_free_mem_cgroup_pages().

In progress of reclaiming, there can be enough memory before try_to_free_xxx is finished. Because we throttle the number of reclaimers, it's better to wake up waiters if there is enough room, in moderate way.
This decreases the system idle time under memory pressure in cgroup.

Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

mm/memcontrol.c | 7 +++++++
1 file changed, 7 insertions(+)

Index: linux-2.6.24-rc3-mm2/mm/memcontrol.c

```
=====
--- linux-2.6.24-rc3-mm2.orig/mm/memcontrol.c
+++ linux-2.6.24-rc3-mm2/mm/memcontrol.c
@@ -816,6 +816,13 @@ void mem_cgroup_uncharge(struct page_cgr
    __mem_cgroup_remove_list(pc);
    spin_unlock_irqrestore(&mz->lru_lock, flags);
    kfree(pc);
+ /*
+  * If there is enough room but there are waiters,
+  * wake up one. (wake up all is tend to be heavy)
+  */
+ if (!res_counter_above_high_watermark(&mem->res) &&
+     waitqueue_active(&mem->waitq))
+   wake_up(&mem->waitq);
+ }
+ }
+ }
```

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