

---

Subject: Re: [RFC][PATCH] memory cgroup enhancements updated [1/10]  
try\_to\_free\_mem\_cgroup\_pages bugfix  
Posted by [Balbir Singh](#) on Tue, 23 Oct 2007 04:00:53 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

KAMEZAWA Hiroyuki wrote:

> Because NODE\_DATA(node)->node\_zonelists[] is guaranteed to contain  
> all necessary zones, it is not necessary to use for\_each\_online\_node.  
>  
> And this for\_each\_online\_node() makes reclaim routine start always  
> from node 0. This is bad.  
>

But with this change, we'll start reclaim on the node that the task  
hit the limit on - right?

> Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>  
>  
>  
>  
> mm/vmscan.c | 8 +++-----  
> 1 file changed, 3 insertions(+), 5 deletions(-)  
>  
> Index: devel-2.6.23-mm1/mm/vmscan.c  
> =====  
> --- devel-2.6.23-mm1.orig/mm/vmscan.c  
> +++ devel-2.6.23-mm1/mm/vmscan.c  
> @@ -1375,15 +1375,13 @@ unsigned long try\_to\_free\_mem\_cgroup\_pag  
> .mem\_cgroup = mem\_cont,  
> .isolate\_pages = mem\_cgroup\_isolate\_pages,  
> };  
> - int node;  
> + int node = numa\_node\_id();  
> struct zone \*\*zones;  
> int target\_zone = gfp\_zone(GFP\_HIGHUSER\_MOVABLE);  
>  
> - for\_each\_online\_node(node) {  
> - zones = NODE\_DATA(node)->node\_zonelists[target\_zone].zones;  
> - if (do\_try\_to\_free\_pages(zones, sc.gfp\_mask, &sc))  
> + zones = NODE\_DATA(node)->node\_zonelists[target\_zone].zones;  
> + if (do\_try\_to\_free\_pages(zones, sc.gfp\_mask, &sc))  
> return 1;  
> - }  
> return 0;  
> }  
> #endif  
>

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---

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

---