
Subject: Re: [RFC][for -mm] memory controller enhancements for NUMA [1/10]
record nid/zid on page_cgroup
Posted by [yamamoto](#) on Wed, 14 Nov 2007 09:22:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

> This patch adds nid/zoneid value to page cgroup.
> This helps per-zone accounting for memory cgroup and reclaim routine.
>
> Signed-off-by: KAMEZAWA Hiroyuki <kmaezawa.hiroyu@jp.fujitsu.com>
>
> mm/memcontrol.c | 6 ++++++
> 1 file changed, 6 insertions(+)
>
> Index: linux-2.6.24-rc2-mm1/mm/memcontrol.c
> ======
> --- linux-2.6.24-rc2-mm1.orig/mm/memcontrol.c
> +++ linux-2.6.24-rc2-mm1/mm/memcontrol.c
> @@ -131,6 +131,8 @@ struct page_cgroup {
> atomic_t ref_cnt; /* Helpful when pages move b/w */
> /* mapped and cached states */
> int flags;
> + short nid;
> + short zid;
> };
> #define PAGE_CGROUP_FLAG_CACHE (0x1) /* charged as cache */
> #define PAGE_CGROUP_FLAG_ACTIVE (0x2) /* page is active in this cgroup */
> @@ -216,6 +218,10 @@ void page_assign_page_cgroup(struct page
> VM_BUG_ON(!page_cgroup_locked(page));
> locked = (page->page_cgroup & PAGE_CGROUP_LOCK);
> page->page_cgroup = ((unsigned long)pc | locked);
> + if (pc) {
> + pc->nid = page_to_nid(page);
> + pc->zid = page_zonenum(page);
> + }
> }
>
> struct page_cgroup *page_get_page_cgroup(struct page *page)

are they worth to be cached?
can't you use page_zonenum(pc->page)?

YAMAMOTO Takashi

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