
Subject: Re: [RFC][PATCH] memory controller per zone patches take 2 [3/10] add per zone active/inactive count

Posted by [Balbir Singh](#) on Sat, 17 Nov 2007 16:07:17 GMT

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

KAMEZAWA Hiroyuki wrote:

- > Counting active/inactive per-zone in memory controller.
- >
- > This patch adds per-zone status in memory cgroup.
- > These values are often read (as per-zone value) by page reclaiming.
- >
- > In current design, per-zone stat is just a unsigned long value and
- > not an atomic value because they are modified only under lru_lock.
- > (for avoiding atomic_t ops.)
- >
- > This patch adds ACTIVE and INACTIVE per-zone status values.
- >
- > For handling per-zone status, this patch adds
- > struct mem_cgroup_per_zone {
- > ...
- > }
- > and some helper functions. This will be useful to add per-zone objects
- > in mem_cgroup.
- >
- > This patch turns memory controller's early_init to be 0 for calling
- > kmalloc().
- >
- > Changelog V1 -> V2
- > - added mem_cgroup_per_zone struct.
- > This will help following patches to implement per-zone objects and
- > pack them into a struct.
- > - added __mem_cgroup_add_list() and __mem_cgroup_remove_list()
- > - fixed page migration handling.
- > - renamed zstat to info (per-zone-info)
- > This will be place for per-zone information(lru, lock, ..)
- > - use page_cgroup_nid()/zid() funcs.
- >
- > Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

The code looks OK to me, pending test on a real NUMA box

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

--

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>
