Subject: Re: [RFC][PATCH] memory controller per zone patches take 2 [8/10] changes in vmscan.c
Posted by Balbir Singh on Sat, 17 Nov 2007 17:37:46 GMT
View Forum Message <> Reply to Message

KAMEZAWA Hiroyuki wrote:
> When using memory controller, there are 2 levels of memory reclaim.
> 1. zone memory reclaim because of system/zone memory shortage.
> 2. memory cgroup memory reclaim because of hitting limit.
>
> These two can be distinguished by sc->mem_cgroup parameter.
>
> This patch tries to make memory cgroup reclaim routine avoid affecting
> system/zone memory reclaim. This patch inserts if (!sc->mem_cgroup) and
> hook to memory_cgroup reclaim support functions.
>
> This patch can be a help for isolating system lru activity and group lru
> activity and shows what additional functions are necessary.
>
> * mem_cgroup_calc_mapped_ratio() ... calculate mapped ratio for cgroup.
> * mem_cgroup_reclaim_imbalance() ... calculate active/inactive balance in
>   cgroup.
> * mem_cgroup_calc_reclaim_active() ... calculate the number of active pages to
>   be scanned in this priority in mem_cgroup.
>
> * mem_cgroup_calc_reclaim_inactive() ... calculate the number of inactive pages
>   to be scanned in this priority in mem_cgroup.
>
> * mem_cgroup_all_unreclaimable() .. checks cgroup's page is all unreclaimable
>   or not.
> * mem_cgroup_get_reclaim_priority() ...
> * mem_cgroup_note_reclaim_priority() ... record reclaim priority (temporal)
> * mem_cgroup_remember_reclaim_priority()
>   .... record reclaim priority as
>   zone->prev_priority.
>   This value is used for calc reclaim_mapped.
> Changelog:
> - merged calc_reclaim_mapped patch in previous version.
>
> Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

The overall idea looks good, it brings the two reclaims closer. The one
pending to do for memory controllers is to make the reclaim lumpy
reclaim aware. But at this point, I don't see a need for it, since
we track only order 1 allocations in the memory controller.

--
 Warm Regards,

Balbir Singh
Linux Technology Center
IBM, ISTL