
Subject: Re: [RFC][PATCH] memory controller per zone patches take 2 [4/10]
calculate mapped ratio for memory
Posted by [Balbir Singh](#) on Sat, 17 Nov 2007 16:12:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

KAMEZAWA Hiroyuki wrote:

> Define function for calculating mapped_ratio in memory cgroup.
>

Could you explain what the ratio is used for? Is it for reclaim
later?

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

>
> include/linux/memcontrol.h | 11 +++++++
> mm/memcontrol.c | 13 +++++++
> 2 files changed, 23 insertions(+), 1 deletion(-)
>
> 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
> @@ -423,6 +423,19 @@ void mem_cgroup_move_lists(struct page_c
> spin_unlock(&mem->lru_lock);
> }
>
> /*
> + * Calculate mapped_ratio under memory controller.
> + */
> +int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem)
> +{
> + s64 total, rss;
> +
> + /* usage is recorded in bytes */
> + total = mem->res.usage >> PAGE_SHIFT;
> + rss = mem_cgroup_read_stat(&mem->stat, MEM_CGROUP_STAT_RSS);
> + return (rss * 100) / total;

Never tried 64 bit division on a 32 bit system. I hope we don't
have to resort to do_div() sort of functionality.

> +}
> +
> + unsigned long mem_cgroup_isolate_pages(unsigned long nr_to_scan,
> + struct list_head *dst,

```
>     unsigned long *scanned, int order,
> Index: linux-2.6.24-rc2-mm1/include/linux/memcontrol.h
> =====
> --- linux-2.6.24-rc2-mm1.orig/include/linux/memcontrol.h
> +++ linux-2.6.24-rc2-mm1/include/linux/memcontrol.h
> @@ -61,6 +61,12 @@ extern int mem_cgroup_prepare_migration(
> extern void mem_cgroup_end_migration(struct page *page);
> extern void mem_cgroup_page_migration(struct page *page, struct page *newpage);
>
> +/*
> + * For memory reclaim.
> + */
> +extern int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem);
> +
> +
> #else /* CONFIG_CGROUP_MEM_CONT */
> static inline void mm_init_cgroup(struct mm_struct *mm,
>     struct task_struct *p)
> @@ -132,7 +138,10 @@ mem_cgroup_page_migration(struct page *p
> {
> }
>
> -
> +static inline int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem)
> +{
> +    return 0;
> +}
> #endif /* CONFIG_CGROUP_MEM_CONT */
>
> #endif /* _LINUX_MEMCONTROL_H */
>
```

--
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>
