
Subject: Re: [RFC][PATCH] memory controller per zone patches take 2 [4/10]
calculate mapped ratio for memory

Posted by [yamamoto](#) on Thu, 22 Nov 2007 08:46:47 GMT

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

> On Thu, 22 Nov 2007 17:34:20 +0900 (JST)

> yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

```
>
> > > > + /* 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.

> > > >

> > > Hmm, maybe it's better to make these numebrs be just "long".

> > > I'll try to change per-cpu-counter implementation.

> > >

> > > Thanks,

> > > -Kame

> >

> > besides that, i think 'total' can be zero here.

> >

> Ah, This is what I do now.

> ==

> +/*

> + * Calculate mapped_ratio under memory controller. This will be used in

> + * vmscan.c for deteremining we have to reclaim mapped pages.

> + */

> +int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem)

> +{

> + long total, rss;

> +

> + /*

> + * usage is recorded in bytes. But, here, we assume the number of

> + * physical pages can be represented by "long" on any arch.

> + */

> + total = (long) (mem->res.usage >> PAGE_SHIFT);

> + rss = (long)mem_cgroup_read_stat(&mem->stat, MEM_CGROUP_STAT_RSS);

> + return (int)((rss * 100L) / total);

> +}

> ==

>

> maybe works well.

>

> -Kame

i meant that "/" total" can cause a division-by-zero exception.

YAMAMOTO Takashi

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
