

hi,

thanks for comments.

> Hi,

>

> On Fri, 7 Sep 2007 12:39:42 +0900 (JST)

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

>

> > +enum mem\_container\_stat\_index {

> > + /\*

> > + \* for MEM\_CONTAINER\_TYPE\_ALL, usage == pagecache + rss

> > + \*/

> > + MEMCONT\_STAT\_PAGECACHE,

> > + MEMCONT\_STAT\_RSS,

> > +

> > + /\*

> > + \* redundant; usage == charge - uncharge

> > + \*/

> > + MEMCONT\_STAT\_CHARGE,

> > + MEMCONT\_STAT\_UNCHARGE,

> > +

> > + /\*

> > + \* mostly for debug

> > + \*/

> > + MEMCONT\_STAT\_ISOLATE,

> > + MEMCONT\_STAT\_ISOLATE\_FAIL,

> > + MEMCONT\_STAT\_NSTATS,

> > +};

> > +

> please add comments on each statistics name.

sure.

> It's uneasy to catch the meaning of

> ISOLATE/ISOLATE\_FAIL without comments.

they aren't useful for users who don't read the relevant code.

probably they should be just removed.

> > +static const char \* const mem\_container\_stat\_desc[] = {

> > + [MEMCONT\_STAT\_PAGECACHE] = "page\_cache",

> > + [MEMCONT\_STAT\_RSS] = "rss",

> > + [MEMCONT\_STAT\_CHARGE] = "charge",

```

> > + [MEMCONT_STAT_UNCHARGE] = "uncharge",
> > + [MEMCONT_STAT_ISOLATE] = "isolate",
> > + [MEMCONT_STAT_ISOLATE_FAIL] = "isolate_fail",
> > +};
> > +
> > +struct mem_container_stat {
> > + atomic_t count[MEMCONT_STAT_NSTATS];
> > +};
> > +
> > +static void mem_container_stat_inc(struct mem_container_stat * stat,
> > + enum mem_container_stat_index idx)
> > +{
> > +
> > + atomic_inc(&stat->count[idx]);
> > +}
> > +
> > +static void mem_container_stat_dec(struct mem_container_stat * stat,
> > + enum mem_container_stat_index idx)
> > +{
> > +
> > + atomic_dec(&stat->count[idx]);
> > +}
> > +
>
> Can we do this accounting as mod_zone_page_state()(in mm/vmstat.c) ?
> (use per-cpu data for accounting.)

```

we can do so later.

```

> > +/* XXX hack; shouldn't be here. it really belongs to struct page_container. */
> > +#define PAGE_CONTAINER_CACHE_BIT 0x1
> > +#define PAGE_CONTAINER_CACHE (1 << PAGE_CONTAINER_CACHE_BIT)
> > +
>
> Is this used for remembering whether a page is charged as page-cache or not ?

```

yes.

```

> > + page_assign_page_container_flags(page,
> > + is_cache ? PAGE_CONTAINER_CACHE : 0, pc);
> > +
> > + stat = &mem->stat;
> > + if (is_cache) {
> > + mem_container_stat_inc(stat, MEMCONT_STAT_PAGECACHE);
> > + } else {
> > + mem_container_stat_inc(stat, MEMCONT_STAT_RSS);
> > + }
>

```

> nitpick,in linux style, one-sentence block shouldn't have braces {}.

>

> ==

> if (is\_cache)

> mem\_cont...

> else

> mem\_cont...

> ==

sure.

> > + mem\_container\_stat\_inc(stat, MEMCONT\_STAT\_CHARGE);

> >

> > spin\_lock\_irqsave(&mem->lru\_lock, flags);

> > list\_add(&pc->lru, &mem->active\_list);

> > @@ -377,6 +454,12 @@ err:

> > return -ENOMEM;

> > }

> >

> > +int mem\_container\_charge(struct page \*page, struct mm\_struct \*mm)

> > +{

> > +

> > + return mem\_container\_charge\_common(page, mm, 0);

> > +}

> > +

> > /\*

> > \* See if the cached pages should be charged at all?

> > \*/

> > @@ -388,7 +471,7 @@ int mem\_container\_cache\_charge(struct pa

> >

> > mem = rcu\_dereference(mm->mem\_container);

> > if (mem->control\_type == MEM\_CONTAINER\_TYPE\_ALL)

> > - return mem\_container\_charge(page, mm);

> > + return mem\_container\_charge\_common(page, mm, 1);

> > else

> > return 0;

> > }

> > @@ -411,15 +494,29 @@ void mem\_container\_uncharge(struct page\_

> > return;

> >

> > if (atomic\_dec\_and\_test(&pc->ref\_cnt)) {

> > + struct mem\_container\_stat \*stat;

> > + int is\_cache;

> > +

> > page = pc->page;

> > lock\_page\_container(page);

> > mem = pc->mem\_container;

> > css\_put(&mem->css);

> > + /\* XXX \*/  
> This kind of comment is bad.

sure.

```
> > + is_cache = (page->page_container & PAGE_CONTAINER_CACHE) != 0;
> >   page_assign_page_container(page, NULL);
> >   unlock_page_container(page);
> >   res_counter_uncharge(&mem->res, 1);
> >
> > + stat = &mem->stat;
> > + if (is_cache) {
> > +   mem_container_stat_dec(stat, MEMCONT_STAT_PAGECACHE);
> > + } else {
> > +   mem_container_stat_dec(stat, MEMCONT_STAT_RSS);
> > + }
> > + mem_container_stat_inc(stat, MEMCONT_STAT_UNCHARGE);
> > +
> >   spin_lock_irqsave(&mem->lru_lock, flags);
> > + BUG_ON(list_empty(&pc->lru));
>
> Why this BUG_ON() is added ?
>
> Thanks
> -Kame
```

to ensure that my understanding is correct.

YAMAMOTO Takashi

---

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

---