
Subject: Re: [PATCH v4 06/11] memcg: add dirty page accounting infrastructure
Posted by [Wu Fengguang](#) on Fri, 29 Oct 2010 11:13:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, Oct 29, 2010 at 03:09:09PM +0800, Greg Thelen wrote:

```
> +
> + case MEMCG_NR_FILE_DIRTY:
> + /* Use Test{Set,Clear} to only un/charge the memcg once. */
> + if (val > 0) {
> +   if (TestSetPageCgroupFileDirty(pc))
> +     val = 0;
> + } else {
> +   if (!TestClearPageCgroupFileDirty(pc))
> +     val = 0;
> + }
```

I'm wondering why TestSet/TestClear and even the cgroup page flags for dirty/writeback/unstable pages are necessary at all (it helps to document in changelog if there are any). For example, VFS will call TestSetPageDirty() before calling mem_cgroup_inc_page_stat(MEMCG_NR_FILE_DIRTY), so there should be no chance of false double counting.

Thanks,
Fengguang

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