
Subject: [PATCH][for -mm] per-zone and reclaim enhancements for memory controller take 3 [4/10] calculate map

Posted by [KAMEZAWA Hiroyuki](#) on Tue, 27 Nov 2007 03:01:15 GMT

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

Define function for calculating mapped_ratio in memory cgroup.

Changelog V1->V2

- Fixed possible divide-by-zero bug.
- Use "long" to avoid 64bit division on 32 bit system.
and does necessary type casts.
- Added comments.

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

```
include/linux/memcontrol.h | 11 ++++++++  
mm/memcontrol.c            | 17 ++++++++  
2 files changed, 27 insertions(+), 1 deletion(-)
```

Index: linux-2.6.24-rc3-mm1/mm/memcontrol.c

```
=====
```

--- linux-2.6.24-rc3-mm1.orig/mm/memcontrol.c 2007-11-26 16:39:02.000000000 +0900
+++ linux-2.6.24-rc3-mm1/mm/memcontrol.c 2007-11-26 16:41:34.000000000 +0900

@ @ -421,6 +421,23 @ @

```
    spin_unlock(&mem->lru_lock);  
}
```

+/*

+ * Calculate mapped_ratio under memory controller. This will be used in
+ * vmscan.c for determining 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) + 1L;

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

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

+}

+

```
    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-rc3-mm1/include/linux/memcontrol.h

```

=====
--- linux-2.6.24-rc3-mm1.orig/include/linux/memcontrol.h 2007-11-26 15:31:19.000000000 +0900
+++ linux-2.6.24-rc3-mm1/include/linux/memcontrol.h 2007-11-26 16:39:05.000000000 +0900
@@ -61,6 +61,12 @@
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 @@
{
}

-
+static inline int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem)
+{
+ return 0;
+}
#endif /* CONFIG_CGROUP_MEM_CONT */

#endif /* LINUX_MEMCONTROL_H */

```