
Subject: [RSS Controller -v2] Fix freeing of active pages
Posted by [Balbir Singh](#) on Thu, 26 Apr 2007 09:28:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

This patch fixes the bad_page() warning seen while freeing a container page. By default all container pages are added to the active list on the container. This patch lazily moves pages to the right list, so that a page on the active list of the zone LRU and the inactive list of the container does not get freed. If that happens, we see a warning from bad_page().

Signed-off-by: Balbir Singh <balbir@linux.vnet.ibm.com>

```
linux/rss_container.h |  0
mm/rss_container.c   | 30 ++++++-----+
2 files changed, 22 insertions(+), 8 deletions(-)
```

```
diff -puN mm/rss_container.c~rss-fix-free-of-active-pages mm/rss_container.c
--- linux-2.6.20/mm/rss_container.c~rss-fix-free-of-active-pages 2007-04-25 22:05:54.000000000
+0530
+++ linux-2.6.20-balbir/mm/rss_container.c 2007-04-25 23:12:35.000000000 +0530
@@ -119,18 +119,36 @@ void container_rss_move_lists(struct pag
 }
```

```
static unsigned long isolate_container_pages(unsigned long nr_to_scan,
- struct list_head *src, struct list_head *dst,
- unsigned long *scanned, struct zone *zone)
+ struct rss_container *rss, struct list_head *dst,
+ unsigned long *scanned, struct zone *zone, bool active)
{
    unsigned long nr_taken = 0;
    struct page *page;
    struct page_container *pc;
    unsigned long scan;
+ struct list_head *src;
    LIST_HEAD(pc_list);

+ src = active ? &rss->active_list : &rss->inactive_list;
+
    for (scan = 0; scan < nr_to_scan && !list_empty(src); scan++) {
        pc = list_entry(src->prev, struct page_container, list);
        page = pc->page;
+
+ /*
+ * We might have got our active, inactive lists
+ * incorrect, fix it here
+ */
+ if (active && !PageActive(page)) {
```

```

+ list_move(&pc->list, &rss->inactive_list);
+ scan--;
+ continue;
+ } else if (!active && PageActive(page)) {
+ list_move(&pc->list, &rss->active_list);
+ scan--;
+ continue;
+ }
+
if (page_zone(page) != zone)
    continue;

@@ -158,12 +176,8 @@ unsigned long isolate_pages_in_container
unsigned long ret;

spin_lock(&rss->res.lock);
- if (active)
- ret = isolate_container_pages(nr_to_scan, &rss->active_list,
- dst, scanned, zone);
- else
- ret = isolate_container_pages(nr_to_scan, &rss->inactive_list,
- dst, scanned, zone);
+ ret = isolate_container_pages(nr_to_scan, rss, dst, scanned, zone,
+ active);
spin_unlock(&rss->res.lock);
return ret;
}
diff -puN include/linux/rss_container.h~rss-fix-free-of-active-pages include/linux/rss_container.h
-
```

--
Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
