Subject: [RFC][PATCH] memory cgroup enhancements updated [7/10] RSS/CACHE failcnt
Posted by KAMEZAWA Hiroyuki on Fri, 19 Oct 2007 09:34:52 GMT
View Forum Message <> Reply to Message

cgroup's resource has "failure" counter. But I think memory cgroup
has 2 types of failure
 - failure of cache
 - failure of RSS

This patch adds above 2 information to stat file.
Above information is shown in "byte". But I wonder showing just counter
is better o not...rather than PAGE_SIZE.

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

 mm/memcontrol.c | 13 ++++++++++++
 1 file changed, 13 insertions(+)

Index: devel-2.6.23-mm1/mm/memcontrol.c
===
--- devel-2.6.23-mm1.orig/mm/memcontrol.c
+++ devel-2.6.23-mm1/mm/memcontrol.c
@@ -50,6 +50,11 @@ enum mem_cgroup_stat_index {
   */
  MEM_CGROUP_STAT_ACTIVE,   /* # of pages in active list */
  MEM_CGROUP_STAT_INACTIVE, /* # of pages on inactive list */
+ /*
+ * precise failcnt
+ */
+ MEM_CGROUP_STAT_FAIL_RSS, /* # of failure in RSS charging */
+ MEM_CGROUP_STAT_FAIL_CACHE,/* # of failure in CACHE charging */

  MEM_CGROUP_STAT_NSTATS,
 };
@@ -486,6 +491,12 @@ retry:
   */
  while (res_counter_charge(&mem->res, PAGE_SIZE)) {
   bool is_atomic = gfp_mask & GFP_ATOMIC;
+ if (is_cache)
+   mem_cgroup_stat_inc(&mem->stat,
+   MEM_CGROUP_STAT_FAIL_CACHE);
+ else
+   mem_cgroup_stat_inc(&mem->stat,
+   MEM_CGROUP_STAT_FAIL_RSS);
   /*
    * We cannot reclaim under GFP_ATOMIC, fail the charge

```
    */
@@ -850,6 +861,8 @@ static const struct mem_cgroup_stat_desc
  [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
  [MEM_CGROUP_STAT_ACTIVE] = { "active", PAGE_SIZE, },
  [MEM_CGROUP_STAT_INACTIVE] = { "inactive", PAGE_SIZE, },
+ [MEM_CGROUP_STAT_FAIL_RSS] = { "rss_failure", PAGE_SIZE, },
+ [MEM_CGROUP_STAT_FAIL_CACHE] = { "cache_failure", PAGE_SIZE, },
 };

 static int mem_control_stat_show(struct seq_file *m, void *arg)
```

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