

---

Subject: [RFC][ for -mm] memory controller enhancements for NUMA [5/10] calc active/inactive impabance

Posted by [KAMEZAWA Hiroyuki](#) on Wed, 14 Nov 2007 08:48:28 GMT

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

---

Define function for determining active/inactive balance in memory cgroup. This imbalance value is calculated against the whole zone..

maybe should be per-zone...?

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

include/linux/memcontrol.h | 8 +++++++  
mm/memcontrol.c | 16 ++++++++  
2 files changed, 24 insertions(+)

Index: linux-2.6.24-rc2-mm1/mm/memcontrol.c

```
=====
--- linux-2.6.24-rc2-mm1.orig/mm/memcontrol.c
+++ linux-2.6.24-rc2-mm1/mm/memcontrol.c
@@ -412,6 +412,22 @@ int mem_cgroup_calc_mapped_ratio(struct
    rss = mem_cgroup_read_stat(&mem->stat, MEM_CGROUP_STAT_RSS);
    return (rss * 100) / total;
}
+/*
+ * Uses mem_cgroup's imbalance instead of zone's lru imbalance.
+ * This will be used for detemining whether page out routine try to free
+ * mapped pages or not.
+ */
+int mem_cgroup_reclaim_imbalance(struct mem_cgroup *mem)
+{
+ s64 total, active, inactive;
+
+ /* usage is recoreded in bytes */
+ total = mem->res.usage >> PAGE_SHIFT;
+ inactive = mem_cgroup_get_all_zonestat(mem, MEM_CGROUP_ZSTAT_INACTIVE);
+ active = total - inactive;
+
+ return active / (inactive + 1);
+}
```

```
unsigned long mem_cgroup_isolate_pages(unsigned long nr_to_scan,
    struct list_head *dst,
```

Index: linux-2.6.24-rc2-mm1/include/linux/memcontrol.h

```
=====
--- linux-2.6.24-rc2-mm1.orig/include/linux/memcontrol.h
+++ linux-2.6.24-rc2-mm1/include/linux/memcontrol.h
@@ -65,6 +65,8 @@ extern void mem_cgroup_page_migration(st
```

```

* For memory reclaim.
*/
extern int mem_cgroup_calc_mapped_ratio(struct mem_cgroup *mem);
+extern int mem_cgroup_reclaim_imbalance(struct mem_cgroup *mem);
+

#else /* CONFIG_CGROUP_MEM_CONT */
@@ -142,6 +144,12 @@ static inline int mem_cgroup_calc_mapped
{
    return 0;
}
+
+static inline int mem_cgroup_reclaim_imbalance(struct mem_cgroup *mem)
+{
+    return 0;
+}
+
#endif /* CONFIG_CGROUP_MEM_CONT */

#endif /* _LINUX_MEMCONTROL_H */

```

---

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

---