
Subject: [RFC][for -mm] memory controller enhancements for NUMA [3/10] define macro for global lru scan

Posted by KAMEZAWA Hiroyuki on Wed, 14 Nov 2007 08:43:52 GMT

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

add macro scan_global_lru().

This is used to detect which scan_control scans global lru or mem_cgroup lru.

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

```
mm/vmscan.c | 17 ++++++-----  
1 file changed, 12 insertions(+), 5 deletions(-)
```

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

```
=====--- linux-2.6.24-rc2-mm1.orig/mm/vmscan.c  
+++ linux-2.6.24-rc2-mm1/mm/vmscan.c  
@@ -127,6 +127,12 @@ long vm_total_pages; /* The total number  
 static LIST_HEAD(shrinker_list);  
 static DECLARE_RWSEM(shrinker_rwsem);  
  
+#ifdef CONFIG_CGROUP_MEM_CONT  
+#define scan_global_lru(sc) (!sc->mem_cgroup)  
+#else  
+#define scan_global_lru(sc) (1)  
+#endif  
+  
/*  
 * Add a shrinker callback to be called from the vm  
 */  
@@ -1290,11 +1296,12 @@ static unsigned long do_try_to_free_page  
 * Don't shrink slabs when reclaiming memory from  
 * over limit cgroups  
 */  
- if (sc->mem_cgroup == NULL)  
+ if (scan_global_lru(sc)) {  
    shrink_slab(sc->nr_scanned, gfp_mask, lru_pages);  
- if (reclaim_state) {  
-   nr_reclaimed += reclaim_state->reclaimed_slab;  
-   reclaim_state->reclaimed_slab = 0;  
+ if (reclaim_state) {  
+   nr_reclaimed += reclaim_state->reclaimed_slab;  
+   reclaim_state->reclaimed_slab = 0;  
+ }  
}  
total_scanned += sc->nr_scanned;
```

```
if (nr_reclaimed >= sc->swap_cluster_max) {
@@ -1321,7 +1328,7 @@ static unsigned long do_try_to_free_page
    congestion_wait(WRITE, HZ/10);
}
/* top priority shrink_caches still had more to do? don't OOM, then */
- if (!sc->all_unreclaimable && sc->mem_cgroup == NULL)
+ if (!sc->all_unreclaimable && scan_global_lru(sc))
    ret = 1;
out:
/*
```

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