Subject: [-mm PATCH] Memory controller make charging gfp mask aware
Posted by Balbir Singh on Wed, 12 Sep 2007 12:14:40 GMT
View Forum Message <> Reply to Message

Nick Piggin pointed out that swap cache and page cache addition routines
could be called from non GFP_KERNEL contexts. This patch makes the charging
routine aware of the gfp context. Charging might fail if the container is
over it's limit, in which case a suitable error is returned.

This patch was tested on a Powerpc box. I am still looking at being able
to test the path, through which allocations happen in non GFP_KERNEL contexts.

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

 include/linux/memcontrol.h | 12 ++++++++----
 include/linux/swap.h | 3 ++-
 mm/filemap.c | 2 +-
 mm/memcontrol.c | 24 ++++++++++++++++++++++++-------
 mm/memory.c | 10 +++++-----
 mm/migrate.c | 2 +-
 mm/swap_state.c | 2 +-
 mm/swapfile.c | 2 +-
 mm/vmscan.c | 5 +++--
 9 files changed, 39 insertions(+), 23 deletions(-)

diff -puN include/linux/memcontrol.h~memory-controller-make-charging-gfpmask-aware
include/linux/memcontrol.h
---
linux-2.6.23-rc4/include/linux/memcontrol.h~memory-controller-make-charging-gfpmask-aware 20
07-09-11 22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/include/linux/memcontrol.h 2007-09-11 22:23:32.000000000 +0530
@@ -32,7 +32,8 @@ extern void mm_free_container(struct mm_
 extern void page_assign_page_container(struct page *page,
     struct page_container *pc);
 extern struct page_container *page_get_page_container(struct page *page);
-extern int mem_container_charge(struct page *page, struct mm_struct *mm);
+extern int mem_container_charge(struct page *page, struct mm_struct *mm,
+   gfp_t gfp_mask);
 extern void mem_container_uncharge(struct page_container *pc);
 extern void mem_container_move_lists(struct page_container *pc, bool active);
 extern unsigned long mem_container_isolate_pages(unsigned long nr_to_scan,
@@ -42,7 +43,8 @@ extern unsigned long mem_container_isola
     struct mem_container *mem_cont,
     int active);
 extern void mem_container_out_of_memory(struct mem_container *mem);
-extern int mem_container_cache_charge(struct page *page, struct mm_struct *mm);
+extern int mem_container_cache_charge(struct page *page, struct mm_struct *mm,

```
+     gfp_t gfp_mask);
 extern struct mem_container *mm_container(struct mm_struct *mm);

 static inline void mem_container_uncharge_page(struct page *page)
@@ -70,7 +72,8 @@ static inline struct page_container *pag
  return NULL;
 }

-static inline int mem_container_charge(struct page *page, struct mm_struct *mm)
+static inline int mem_container_charge(struct page *page, struct mm_struct *mm,
+     gfp_t gfp_mask)
 {
  return 0;
 }
@@ -89,7 +92,8 @@ static inline void mem_container_move_li
 }

 static inline int mem_container_cache_charge(struct page *page,
-      struct mm_struct *mm)
+       struct mm_struct *mm,
+       gfp_t gfp_mask)
 {
  return 0;
 }
diff -puN mm/memcontrol.c~memory-controller-make-charging-gfpmask-aware mm/memcontrol.c
---
linux-2.6.23-rc4/mm/memcontrol.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/memcontrol.c 2007-09-12 00:25:12.000000000 +0530
@@ -261,7 +261,8 @@ unsigned long mem_container_isolate_page
  * 0 if the charge was successful
  * < 0 if the container is over its limit
  */
-int mem_container_charge(struct page *page, struct mm_struct *mm)
+int mem_container_charge(struct page *page, struct mm_struct *mm,
+    gfp_t gfp_mask)
 {
  struct mem_container *mem;
  struct page_container *pc, *race_pc;
@@ -287,7 +288,7 @@ int mem_container_charge(struct page *pa

  unlock_page_container(page);

- pc = kzalloc(sizeof(struct page_container), GFP_KERNEL);
+ pc = kzalloc(sizeof(struct page_container), gfp_mask);
  if (pc == NULL)
   goto err;
```

```
@@ -314,7 +315,14 @@ int mem_container_charge(struct page *pa
   * the container limit.
   */
  while (res_counter_charge(&mem->res, 1)) {
-  if (try_to_free_mem_container_pages(mem))
+  bool is_atomic = gfp_mask & GFP_ATOMIC;
+  /*
+   * We cannot reclaim under GFP_ATOMIC, fail the charge
+   */
+  if (is_atomic)
+   goto noreclaim;
+
+  if (try_to_free_mem_container_pages(mem, gfp_mask))
    continue;

   /*
@@ -338,9 +346,10 @@ int mem_container_charge(struct page *pa
    congestion_wait(WRITE, HZ/10);
    continue;
   }
-
+noreclaim:
   css_put(&mem->css);
-  mem_container_out_of_memory(mem);
+  if (!is_atomic)
+   mem_container_out_of_memory(mem);
   goto free_pc;
  }

@@ -381,7 +390,8 @@ err:
 /*
  * See if the cached pages should be charged at all?
  */
-int mem_container_cache_charge(struct page *page, struct mm_struct *mm)
+int mem_container_cache_charge(struct page *page, struct mm_struct *mm,
+    gfp_t gfp_mask)
 {
  struct mem_container *mem;
  if (!mm)
@@ -389,7 +399,7 @@ int mem_container_cache_charge(struct pa

  mem = rcu_dereference(mm->mem_container);
  if (mem->control_type == MEM_CONTAINER_TYPE_ALL)
-  return mem_container_charge(page, mm);
+  return mem_container_charge(page, mm, gfp_mask);
  else
   return 0;
 }
```

```
diff -puN mm/memory.c~memory-controller-make-charging-gfpmask-aware mm/memory.c
--- linux-2.6.23-rc4/mm/memory.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/memory.c 2007-09-11 22:54:09.000000000 +0530
@@ -1137,7 +1137,7 @@ static int insert_page(struct mm_struct
  pte_t *pte;
  spinlock_t *ptl;

- retval = mem_container_charge(page, mm);
+ retval = mem_container_charge(page, mm, GFP_KERNEL);
  if (retval)
   goto out;

@@ -1638,7 +1638,7 @@ gotten:
   goto oom;
  cow_user_page(new_page, old_page, address, vma);

- if (mem_container_charge(new_page, mm))
+ if (mem_container_charge(new_page, mm, GFP_KERNEL))
   goto oom_free_new;

  /*
@@ -2101,7 +2101,7 @@ static int do_swap_page(struct mm_struct
  }

  delayacct_clear_flag(DELAYACCT_PF_SWAPIN);
- if (mem_container_charge(page, mm)) {
+ if (mem_container_charge(page, mm, GFP_KERNEL)) {
   ret = VM_FAULT_OOM;
   goto out;
  }
@@ -2185,7 +2185,7 @@ static int do_anonymous_page(struct mm_s
  if (!page)
   goto oom;

- if (mem_container_charge(page, mm))
+ if (mem_container_charge(page, mm, GFP_KERNEL))
   goto oom_free_page;

  entry = mk_pte(page, vma->vm_page_prot);
@@ -2320,7 +2320,7 @@ static int __do_fault(struct mm_struct *

  }

- if (mem_container_charge(page, mm)) {
+ if (mem_container_charge(page, mm, GFP_KERNEL)) {
   ret = VM_FAULT_OOM;
   goto out;
```

```
 }
diff -puN mm/filemap.c~memory-controller-make-charging-gfpmask-aware mm/filemap.c
--- linux-2.6.23-rc4/mm/filemap.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/filemap.c 2007-09-11 22:54:19.000000000 +0530
@@ -445,7 +445,7 @@ int add_to_page_cache(struct page *page,

  if (error == 0) {

-  error = mem_container_cache_charge(page, current->mm);
+  error = mem_container_cache_charge(page, current->mm, gfp_mask);
   if (error)
    goto out;

diff -puN mm/migrate.c~memory-controller-make-charging-gfpmask-aware mm/migrate.c
--- linux-2.6.23-rc4/mm/migrate.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/migrate.c 2007-09-11 22:54:29.000000000 +0530
@@ -158,7 +158,7 @@ static void remove_migration_pte(struct
    return;
   }

- if (mem_container_charge(new, mm)) {
+ if (mem_container_charge(new, mm, GFP_KERNEL)) {
   pte_unmap(ptep);
   return;
  }
diff -puN mm/page_alloc.c~memory-controller-make-charging-gfpmask-aware mm/page_alloc.c
diff -puN mm/rmap.c~memory-controller-make-charging-gfpmask-aware mm/rmap.c
diff -puN mm/swapfile.c~memory-controller-make-charging-gfpmask-aware mm/swapfile.c
--- linux-2.6.23-rc4/mm/swapfile.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/swapfile.c 2007-09-11 22:54:52.000000000 +0530
@@ -510,7 +510,7 @@ unsigned int count_swap_pages(int type,
 static int unuse_pte(struct vm_area_struct *vma, pte_t *pte,
   unsigned long addr, swp_entry_t entry, struct page *page)
 {
- if (mem_container_charge(page, vma->vm_mm))
+ if (mem_container_charge(page, vma->vm_mm, GFP_KERNEL))
   return -ENOMEM;

  inc_mm_counter(vma->vm_mm, anon_rss);
diff -puN mm/swap_state.c~memory-controller-make-charging-gfpmask-aware mm/swap_state.c
---
linux-2.6.23-rc4/mm/swap_state.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/swap_state.c 2007-09-11 22:55:12.000000000 +0530
@@ -81,7 +81,7 @@ static int __add_to_swap_cache(struct pa
```

```
  error = radix_tree_preload(gfp_mask);
  if (!error) {

-  error = mem_container_cache_charge(page, current->mm);
+  error = mem_container_cache_charge(page, current->mm, gfp_mask);
   if (error)
    goto out;

diff -puN mm/vmscan.c~memory-controller-make-charging-gfpmask-aware mm/vmscan.c
--- linux-2.6.23-rc4/mm/vmscan.c~memory-controller-make-charging-gfpmask-aware 2007-09-11
22:20:50.000000000 +0530
+++ linux-2.6.23-rc4-balbir/mm/vmscan.c 2007-09-11 23:05:40.000000000 +0530
@@ -1357,10 +1357,11 @@ unsigned long try_to_free_pages(struct z
 #define ZONE_USERPAGES ZONE_NORMAL
 #endif

-unsigned long try_to_free_mem_container_pages(struct mem_container *mem_cont)
+unsigned long try_to_free_mem_container_pages(struct mem_container *mem_cont,
+      gfp_t gfp_mask)
 {
  struct scan_control sc = {
-  .gfp_mask = GFP_KERNEL,
+  .gfp_mask = gfp_mask,
   .may_writepage = !laptop_mode,
   .may_swap = 1,
   .swap_cluster_max = SWAP_CLUSTER_MAX,
diff -puN include/linux/swap.h~memory-controller-make-charging-gfpmask-aware
include/linux/swap.h
---
linux-2.6.23-rc4/include/linux/swap.h~memory-controller-make-charging-gfpmask-aware 2007-09-
12 00:11:37.000000000 +0530
+++ linux-2.6.23-rc4-balbir/include/linux/swap.h 2007-09-11 23:05:59.000000000 +0530
@@ -191,7 +191,8 @@ extern void swap_setup(void);
 /* linux/mm/vmscan.c */
 extern unsigned long try_to_free_pages(struct zone **zones, int order,
      gfp_t gfp_mask);
-extern unsigned long try_to_free_mem_container_pages(struct mem_container *mem);
+extern unsigned long try_to_free_mem_container_pages(struct mem_container *mem,
+      gfp_t gfp_mask);
 extern int __isolate_lru_page(struct page *page, int mode);
 extern unsigned long shrink_all_memory(unsigned long nr_pages);
 extern int vm_swappiness;

_

--
 Warm Regards,
 Balbir Singh
 Linux Technology Center
```

IBM, ISTL

_____

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