
Subject: [RFC][PATCH 2/3] The accounting hooks and core
Posted by [Pavel Emelianov](#) on Thu, 30 Aug 2007 09:39:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

The struct page gets an extra pointer (just like it has with the RSS controller) and this pointer points to the array of the kmem_container pointers - one for each object stored on that page itself.

Thus the i'th object on the page is accounted to the container pointed by the i'th pointer on that array and when the object is freed we unaccount its size to this particular container, not the container current task belongs to.

This is done so, because the context objects are freed is most often not the same as the one this objects was allocated in (due to RCU and reference counters).

Kmem cache marked as SLAB_CHARGE will perform the accounting.

Signed-off-by: Pavel Emelianov <xemul@openvz.org>

```
include/linux/mm_types.h | 3
include/linux/slab.h      | 1
mm/slub.c                 | 158 ++++++
3 files changed, 161 insertions(+), 1 deletion(-)
```

```
diff --git a/include/linux/mm_types.h b/include/linux/mm_types.h
index 2dd6c53..629499e 100644
--- a/include/linux/mm_types.h
+++ b/include/linux/mm_types.h
@@ -96,6 +96,9 @@ struct page {
    unsigned int gfp_mask;
    unsigned long trace[8];
#endif
+#ifdef CONFIG_CONTAINER_KMEM
+ struct kmem_container **containers;
+#endif
};
```

```
/*
diff --git a/include/linux/slab.h b/include/linux/slab.h
index 3a5bad3..cd7d50d 100644
--- a/include/linux/slab.h
+++ b/include/linux/slab.h
@@ -28,6 +28,7 @@
```

```

#define SLAB_DESTROY_BY_RCU 0x00080000UL /* Defer freeing slabs to RCU */
#define SLAB_MEM_SPREAD 0x00100000UL /* Spread some memory over cpuset */
#define SLAB_TRACE 0x00200000UL /* Trace allocations and frees */
+#define SLAB_CHARGE 0x00400000UL /* Charge allocations */

/* The following flags affect the page allocator grouping pages by mobility */
#define SLAB_RECLAIM_ACCOUNT 0x00020000UL /* Objects are reclaimable */
diff --git a/mm/slub.c b/mm/slub.c
index 130171d..e4f7390 100644
--- a/mm/slub.c
+++ b/mm/slub.c
@@ -1039,6 +1044,73 @@ static inline void add_full(struct kmem_
static inline void kmem_cache_open_debug_check(struct kmem_cache *s) {}
#define slub_debug 0
#endif
+
+
+#ifdef CONFIG_CONTAINER_KMEM
+/*
+ * Fast path stubs
+ */
+
+static int __kmem_charge(struct kmem_cache *s, void *obj, gfp_t flags);
+static inline
+int kmem_charge(struct kmem_cache *s, void *obj, gfp_t flags)
+{
+ return (s->flags & SLAB_CHARGE) ? __kmem_charge(s, obj, flags) : 0;
+}
+
+static void __kmem_uncharge(struct kmem_cache *s, void *obj);
+static inline
+void kmem_uncharge(struct kmem_cache *s, void *obj)
+{
+ if (s->flags & SLAB_CHARGE)
+ __kmem_uncharge(s, obj);
+}
+
+static int __kmem_prepare(struct kmem_cache *s, struct page *pg, gfp_t flags);
+static inline
+int kmem_prepare(struct kmem_cache *s, struct page *pg, gfp_t flags)
+{
+ return (s->flags & SLAB_CHARGE) ? __kmem_prepare(s, pg, flags) : 0;
+}
+
+static void __kmem_release(struct kmem_cache *s, struct page *pg);
+static inline
+void kmem_release(struct kmem_cache *s, struct page *pg)
+{
+ if (s->flags & SLAB_CHARGE)

```

```

+ __kmem_release(s, pg);
+}
+
+static inline int is_kmalloc_cache(struct kmem_cache *s)
+{
+ int km_idx;
+
+ km_idx = s - kmalloc_caches;
+ return km_idx >= 0 && km_idx < ARRAY_SIZE(kmalloc_caches);
+}
+
+static inline
+int kmem_charge(struct kmem_cache *s, void *obj, gfp_t flags)
+{
+ return 0;
+}
+
+static inline
+void kmem_uncharge(struct kmem_cache *s, void *obj)
+{
+}
+
+static inline
+int kmem_prepare(struct kmem_cache *s, struct page *pg, gfp_t flags)
+{
+ return 0;
+}
+
+static inline
+void kmem_release(struct kmem_cache *s, struct page *pg)
+{
+}
+
+
+/*
+ * Slab allocation and freeing
+ */
@@ -1062,7 +1134,10 @@ static struct page *allocate_slab(struct
    page = alloc_pages_node(node, flags, s->order);

    if (!page)
-   return NULL;
+   goto err_page;
+
+ if (kmem_prepare(s, page, flags) < 0)
+   goto err_prep;

    mod_zone_page_state(page_zone(page),

```

```

(s->flags & SLAB_RECLAIM_ACCOUNT) ?
@@ -1070,6 +1145,11 @@ static struct page *allocate_slab(struct
    pages);

    return page;
+
+err_prep:
+ __free_pages(page, s->order);
+err_page:
+ return NULL;
}

static void setup_object(struct kmem_cache *s, struct page *page,
@@ -1165,6 +1245,8 @@ static void rcu_free_slab(struct rcu_he

static void free_slab(struct kmem_cache *s, struct page *page)
{
+ kmem_release(s, page);
+
    if (unlikely(s->flags & SLAB_DESTROY_BY_RCU)) {
/*
    * RCU free overloads the RCU head over the LRU
@@ -1572,6 +1654,11 @@ static void __always_inline *slab_alloc(
    }
    local_irq_restore(flags);

+ if (object && kmem_charge(s, object, gfpflags) < 0) {
+ kmem_cache_free(s, object);
+ return NULL;
+ }
+
    if (unlikely((gfpflags & __GFP_ZERO) && object))
        memset(object, 0, s->objsize);

@@ -1667,6 +1754,8 @@ static void __always_inline slab_free(st
    void **object = (void *)x;
    unsigned long flags;

+ kmem_uncharge(s, x);
+
    local_irq_save(flags);
    debug_check_no_locks_freed(object, s->objsize);
    if (likely(page == s->cpu_slab[smp_processor_id()]) &&
@@ -3870,6 +4011,78 @@ static int __init slab_sysfs_init(void)
    return css_to_kmem(task_subsys_state(tsk, kmem_subsys_id));
}

+static int __kmem_charge(struct kmem_cache *s, void *obj, gfp_t flags)

```

```

+{
+ struct page *pg;
+ struct kmem_container *cnt;
+ struct kmem_container **obj_container;
+
+ pg = virt_to_head_page(obj);
+ obj_container = pg->containers;
+ if (unlikely(obj_container == NULL)) {
+ /*
+  * turned on after some objects were allocated
+  */
+ if (__kmem_prepare(s, pg, flags) < 0)
+ return -ENOMEM;
+
+ obj_container = pg->containers;
+ }
+
+ cnt = task_kmem_container(current);
+ if (res_counter_charge(&cnt->res, s->size))
+ return -ENOMEM;
+
+ css_get(&cnt->css);
+ obj_container[slab_index(obj, s, page_address(pg))] = cnt;
+ return 0;
+}
+
+static void __kmem_uncharge(struct kmem_cache *s, void *obj)
+{
+ struct page *pg;
+ struct kmem_container *cnt;
+ struct kmem_container **obj_container;
+
+ pg = virt_to_head_page(obj);
+ obj_container = pg->containers;
+ if (obj_container == NULL)
+ return;
+
+ obj_container += slab_index(obj, s, page_address(pg));
+ cnt = *obj_container;
+ if (cnt == NULL)
+ return;
+
+ res_counter_uncharge(&cnt->res, s->size);
+ *obj_container = NULL;
+ css_put(&cnt->css);
+}
+
+static int __kmem_prepare(struct kmem_cache *s, struct page *pg, gfp_t flags)

```

```

+{
+ struct kmem_container **ptr;
+
+ ptr = kzalloc(s->objects * sizeof(struct kmem_container *), flags);
+ if (ptr == NULL)
+ return -ENOMEM;
+
+ pg->containers = ptr;
+ return 0;
+}
+
+static void __kmem_release(struct kmem_cache *s, struct page *pg)
+{
+ struct kmem_container **ptr;
+
+ ptr = pg->containers;
+ if (ptr == NULL)
+ return;
+
+ kfree(ptr);
+ pg->containers = NULL;
+}
+
+/*
+ * containers interface
+ */

```
