
Subject: [PATCH 2/8] Add container pointer on struct page
Posted by [Pavel Emelianov](#) on Mon, 04 Jun 2007 13:26:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Each page is supposed to have an owner - the container that touched the page first. The owner stays alive during the page lifetime even if the task that touched the page dies or moves to another container.

This ownership is the forerunner for the "fair" page sharing accounting, in which page has as many owners as it is really used by.

Page ownership getter/setter are done as static inline function.

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

```
diff -upr linux-2.6.22-rc2-mm1.orig/include/linux/mm.h linux-2.6.22-rc2-mm1-2/include/linux/mm.h
--- linux-2.6.22-rc2-mm1.orig/include/linux/mm.h 2007-05-30 12:32:35.000000000 +0400
+++ linux-2.6.22-rc2-mm1-2/include/linux/mm.h 2007-06-01 16:35:15.000000000 +0400
@@ -249,6 +249,30 @@ struct vm_operations_struct {
```

```
    struct mmu_gather;
    struct inode;
+struct page_container;
+
+#ifdef CONFIG_RSS_CONTAINER
+static inline struct page_container *page_container(struct page *pg)
+{
+ return pg->rss_container;
+}
+
+static inline void set_page_container(struct page *pg,
+ struct page_container *container)
+{
+ pg->rss_container = container;
+}
+#else
+static inline struct page_container *page_container(struct page *pg)
+{
+ return NULL;
+}
+
+static inline void set_page_container(struct page *pg,
+ struct page_container *container)
+{
```

```
+}
+#endif

#define page_private(page) ((page)->private)
#define set_page_private(page, v) ((page)->private = (v))
diff -upr linux-2.6.22-rc2-mm1.orig/include/linux/mm_types.h
linux-2.6.22-rc2-mm1-2/include/linux/mm_types.h
--- linux-2.6.22-rc2-mm1.orig/include/linux/mm_types.h 2007-05-30 12:32:35.000000000 +0400
+++ linux-2.6.22-rc2-mm1-2/include/linux/mm_types.h 2007-06-01 16:35:13.000000000 +0400
@@ -83,6 +83,9 @@ struct page {
    unsigned int gfp_mask;
    unsigned long trace[8];
#endif
+#ifdef CONFIG_RSS_CONTAINER
+ struct page_container *rss_container;
+#endif
};

#endif /* _LINUX_MM_TYPES_H */
```
