

[PATCH 2/9] clockpro-nonresident-del.patch

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-12/msg08494.html>

- *From:* Peter Zijlstra <a.p.zijlstra@xxxxxxxx>
 - *Date:* Fri, 30 Dec 2005 23:42:54 +0100
-

From: Peter Zijlstra <a.p.zijlstra@xxxxxxxx>

Since the coupling of the resident and nonresident clocks depend on the actual number of nonresident pages present, stale entries influence the actual behaviour. Hence we remove the nonresident pages from the map.

Signed-off-by: Peter Zijlstra <a.p.zijlstra@xxxxxxxx>

mm/memory.c | 24 +++++
mm/swapfile.c | 12 ++++++
2 files changed, 34 insertions(+), 2 deletions(-)

Index: linux-2.6-git/mm/swapfile.c

```
-----  
--- linux-2.6-git.orig/mm/swapfile.c  
+++ linux-2.6-git/mm/swapfile.c  
@@ -278,7 +278,8 @@ void swap_free(swp_entry_t entry)  
  
p = swap_info_get(entry);  
if (p) {  
- swap_entry_free(p, swp_offset(entry));  
+ if (!swap_entry_free(p, swp_offset(entry)))  
+ nonresident_get(&swapper_space, entry.val);  
spin_unlock(&swap_lock);  
}  
}  
@@ -375,8 +376,15 @@ void free_swap_and_cache(swp_entry_t ent  
  
p = swap_info_get(entry);  
if (p) {  
- if (swap_entry_free(p, swp_offset(entry)) == 1)  
+ switch (swap_entry_free(p, swp_offset(entry))) {  
+ case 1:  
page = find_trylock_page(&swapper_space, entry.val);  
+ break;  
+  
+ case 0:  
+ nonresident_get(&swapper_space, entry.val);
```

[PATCH 2/9] clockpro-nonresident-del.patch

```
+ break;
+ }
spin_unlock(&swap_lock);
}
if (page) {
Index: linux-2.6-git/mm/memory.c
=====
--- linux-2.6-git.orig/mm/memory.c
+++ linux-2.6-git/mm/memory.c
@@ -595,6 +595,27 @@ int copy_page_range(struct mm_struct *ds
return 0;
}

+static void free_file(struct vm_area_struct *vma,
+ unsigned long offset)
+{
+ struct address_space *mapping;
+ struct page *page;
+
+ if (!vma ||
+ !vma->vm_file ||
+ !vma->vm_file->f_mapping)
+ return;
+
+ mapping = vma->vm_file->f_mapping;
+ page = find_get_page(mapping, offset);
+ if (page) {
+ page_cache_release(page);
+ return;
+ }
+
+ nonresident_get(mapping, offset);
+}
+
static unsigned long zap_pte_range(struct mmu_gather *tlb,
struct vm_area_struct *vma, pmd_t *pmd,
unsigned long addr, unsigned long end,
@@ -610,6 +631,7 @@ static unsigned long zap_pte_range(struc
do {
pte_t ptent = *pte;
if (pte_none(ptent)) {
+ free_file(vma, pte_to_pgoff(ptent));
(*zap_work)--;
continue;
}
@@ -668,6 +690,8 @@ static unsigned long zap_pte_range(struc
continue;
if (!pte_file(ptent))
free_swap_and_cache(pte_to_swp_entry(ptent));
+ else
+ free_file(vma, pte_to_pgoff(ptent));
```

[PATCH 2/9] clockpro-nonresident-del.patch

```
pte_clear_full(mm, addr, pte, tlb->fullmm);  
} while (pte++, addr += PAGE_SIZE, (addr != end && *zap_work > 0));
```

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

• *References:*

- ◆ [\[PATCH\] vm: page-replace and clockpro](#)

◇ *From:* Peter Zijlstra

- Prev by Date: [\[PATCH 09/14\] page-replace-reinsert.patch](#)
- Next by Date: [\[PATCH 04/14\] page-replace-activate page.patch](#)
- Previous by thread: [\[PATCH 09/14\] page-replace-reinsert.patch](#)
- Next by thread: [\[PATCH 04/14\] page-replace-activate page.patch](#)
- Index(es):
 - ◆ [Date](#)
 - ◆ [Thread](#)