

## [PATCH 2/4] 2.6.2-rc2-mm2 CPU Hotplug

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-01/7894.html>

---

**From:** Rusty Russell ([rusty\\_at\\_rustcorp.com.au](mailto:rusty_at_rustcorp.com.au))

**Date:** 01/31/04

To: [akpm@osdl.org](mailto:akpm@osdl.org)

Date: Sun, 01 Feb 2004 01:16:38 +1100

Like CONFIG\_PM, CONFIG\_HOTPLUG\_CPU wants drain\_local\_pages(), except we want to drain another (dead) CPU. Trivial change, mainly fixing up callers.

Name: Change drain\_local\_pages() To Take a CPU Number

Author: Rusty Russell

Status: Booted on 2.6.2-rc2-bk2

D: Hotplug CPU needs to drain pages on a downed CPU (usually it's the

D: current cpu). Make it an argument, and expose it if CONFIG\_HOTPLUG\_CPU

D: as well as CONFIG\_PM.

```
diff -urpN --exclude TAGS -X /home/rusty/devel/kernel/kernel-patches/current-dontdiff --minimal
.6243-linux-2.6.1-bk4/include/linux/suspend.h .6243-linux-2.6.1-bk4.updated/include/linux/suspend.h
--- .6243-linux-2.6.1-bk4/include/linux/suspend.h 2003-10-23 10:02:54.000000000 +1000
+++ .6243-linux-2.6.1-bk4.updated/include/linux/suspend.h 2004-01-19 15:14:01.000000000 +1100
@@ -43,7 +43,7 @@ struct suspend_header {
extern int shrink_mem(void);

/* mm/page_alloc.c */
-extern void drain_local_pages(void);
+extern void drain_local_pages(unsigned int cpu);

extern unsigned int nr_copy_pages __nosavedata;
extern suspend_pagedir_t *pagedir_nosave __nosavedata;
diff -urpN --exclude TAGS -X /home/rusty/devel/kernel/kernel-patches/current-dontdiff --minimal
.6243-linux-2.6.1-bk4/mm/page_alloc.c .6243-linux-2.6.1-bk4.updated/mm/page_alloc.c
--- .6243-linux-2.6.1-bk4/mm/page_alloc.c 2004-01-19 15:13:56.000000000 +1100
+++ .6243-linux-2.6.1-bk4.updated/mm/page_alloc.c 2004-01-19 15:14:01.000000000 +1100
@@ -413,11 +413,13 @@ int is_head_of_free_region(struct page *
spin_unlock_irqrestore(&zone->lock, flags);
return 0;
}
+endif /* CONFIG_PM */

+#if defined(CONFIG_PM) || defined(CONFIG_HOTPLUG_CPU)
/*
```

## Linux-Kernel: [PATCH 2/4] 2.6.2-rc2-mm2 CPU Hotplug

```

* Spill all of this CPU's per-cpu pages back into the buddy allocator.
*/
-void drain_local_pages(void)
+void drain_local_pages(unsigned int cpu)
{
    unsigned long flags;
    struct zone *zone;
@@ -427,7 +429,7 @@ void drain_local_pages(void)
    for_each_zone(zone) {
        struct per_cpu_pageset *pset;

- pset = &zone->pageset[smp_processor_id()];
+ pset = &zone->pageset[cpu];
        for (i = 0; i < ARRAY_SIZE(pset->pcp); i++) {
            struct per_cpu_pages *pcp;

@@ -438,7 +440,7 @@ void drain_local_pages(void)
        }
        local_irq_restore(flags);
    }
-#endif /* CONFIG_PM */
+#endif /* CONFIG_PM || CONFIG_HOTPLUG_CPU */

/*
 * Free a 0-order page
diff -urpN --exclude TAGS -X /home/rusty/devel/kernel/kernel-patches/current-dontdiff --minimal
.6243-linux-2.6.1-bk4/kernel/power/pmdisk.c .6243-linux-2.6.1-bk4.updated/kernel/power/pmdisk.c
--- .6243-linux-2.6.1-bk4/kernel/power/pmdisk.c 2003-10-23 10:02:54.000000000 +1000
+++ .6243-linux-2.6.1-bk4.updated/kernel/power/pmdisk.c 2004-01-19 15:14:01.000000000 +1100
@@ -627,7 +627,7 @@ int pmdisk_suspend(void)
    if ((error = read_swapfiles()))
        return error;

- drain_local_pages();
+ drain_local_pages(smp_processor_id());

    pm_pagedir_nosave = NULL;
    pr_debug("pmdisk: Counting pages to copy.\n" );
@@ -658,7 +658,7 @@ int pmdisk_suspend(void)
    /* During allocating of suspend pagedir, new cold pages may appear.
     * Kill them
     */
- drain_local_pages();
+ drain_local_pages(smp_processor_id());

    /* copy */
    copy_pages();
diff -urpN --exclude TAGS -X /home/rusty/devel/kernel/kernel-patches/current-dontdiff --minimal
.6243-linux-2.6.1-bk4/kernel/power/swsusp.c .6243-linux-2.6.1-bk4.updated/kernel/power/swsusp.c
--- .6243-linux-2.6.1-bk4/kernel/power/swsusp.c 2004-01-10 13:59:39.000000000 +1100
+++ .6243-linux-2.6.1-bk4.updated/kernel/power/swsusp.c 2004-01-19 15:14:01.000000000 +1100

```

## Linux-Kernel: [PATCH 2/4] 2.6.2-rc2-mm2 CPU Hotplug

```
@@ -520,7 +520,7 @@ static int suspend_prepare_image(void)
    struct sysinfo i;
    unsigned int nr_needed_pages = 0;

- drain_local_pages();
+ drain_local_pages(smp_processor_id());

    pagedir_nosave = NULL;
    printk( "/critical section: Counting pages to copy" );
@@ -553,7 +553,7 @@ static int suspend_prepare_image(void)
    nr_copy_pages_check = nr_copy_pages;
    pagedir_order_check = pagedir_order;

- drain_local_pages(); /* During allocating of suspend pagedir, new cold pages may appear. Kill them */
+ drain_local_pages(smp_processor_id()); /* During allocating of suspend pagedir, new cold pages may
appear. Kill them */
    if (nr_copy_pages != count_and_copy_data_pages(pagedir_nosave)) /* copy */
        BUG();

--
    Anyone who quotes me in their sig is an idiot. -- Rusty Russell.
-
To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
the body of a message to majordomo@vger.kernel.org
More majordomo info at http://vger.kernel.org/majordomo-info.html
Please read the FAQ at http://www.tux.org/lkml/
```