

Power Management Update

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-08/8147.html>

From: Patrick Mochel (*mochel_at_osdl.org*)

Date: 08/30/03

Date: Sat, 30 Aug 2003 14:25:07 -0700 (PDT)

To: <linux-kernel@vger.kernel.org>

I'm pleased to announce the release of the first patchset of power management changes for 2.6.0. The purpose of this release is to give people a chance to review and test the PM code before it's sent on to Linus.

These patches include a number of cleanups and fixes to the PM core code, the driver core PM code, and swsusp. I have verified that all suspend states (standby, suspend-to-ram, and suspend-to-disk) work on a number of personal systems using ACPI as the low-level power interface. However, this is with limited functionality (from a VGA console with minimal processes running).

These patches should restore suspend functionality for those that were able to successfully do it before -test3 and -test4. My apologies for the inconvenience my previous changes caused. These patches will probably not allow any more people to suspend/resume than before.

The net benefit of these, and the already committed ones, are a cleaner power management subsystem and the development of the proper framework for successfully suspending and resuming the entire system. There are still several rough edges, though we seem to be making headway on those relatively rapidly, and are my sole focus at the moment.

My main concerns right now are:

- Platform devices, and more generally, devices that may belong to more than one class. It's mainly a driver model problem, though it has PM implications that appear to be holding a few people up.

- Drivers

Drivers have always been the main impedence to having a working PM core, though it's been difficult to make a lot of progress. I have a number of devices that I will verify work properly, and be in contact with the maintainers if necessary. (Though, I seem to be having more problems with IRQ routing at the moment.)

Linux-Kernel: Power Management Update

– Getting it work on more systems.

Hopefully we will not run into any serious issues, though the PM code has traditionally been finicky. I have a wide array of test machines and willing testers, so this should move quickly.

– APM

I unfortunately have not had a chance to look into the reported APM problems. But, I'm happy to say that I finally dug out an old laptop that has APM on it. I should make traction soon.

I encourage willing people to download the patch, test, and report any problems back to me and/or the list. I cannot guarantee definite or timely results for systems where PM simply doesn't work. However, the more systems we characterize, the easier this will become in the future. Please be patient.

If you're using BitKeeper, you can pull the tree from:

```
bk://kernel.bkbits.net:/home/mochel/linux-2.5-power
```

Or, a GNU patch is available at:

<http://developer.osdl.org/~mochel/patches/test4-pm1/test4-pm1.diff.bz2>

There are split patches for each BK revision in that directory. The changelogs are appended.

Pat

This will update the following files:

```
arch/i386/kernel/suspend.c | 141 -----
arch/i386/kernel/suspend_asm.S | 94 -----
arch/i386/Makefile | 1
arch/i386/kernel/Makefile | 2
arch/i386/power/Makefile | 2
arch/i386/power/cpu.c | 141 ++++++++
arch/i386/power/swsusp.S | 104 ++++++
drivers/acpi/sleep/main.c | 53 +---
drivers/acpi/sleep/proc.c | 73 ++++++
drivers/acpi/sleep/sleep.h | 3
drivers/base/core.c | 33 +---
drivers/base/power/main.c | 13 -
drivers/base/power/power.h | 3
drivers/base/power/resume.c | 21 +
drivers/base/power/suspend.c | 10
include/asm-i386/suspend.h | 7
include/linux/suspend.h | 1
kernel/power/Makefile | 2
kernel/power/console.c | 2
kernel/power/disk.c | 337 ++++++
```

Linux-Kernel: Power Management Update

```
kernel/power/main.c | 450 ++++++-----  
kernel/power/power.h | 39 +---  
kernel/power/swsusp.c | 357 ++++++-----  
kernel/sys.c | 2  
24 files changed, 1053 insertions(+), 838 deletions(-)
```

through these ChangeSets:

<mochel@osdl.org> (03/08/30 1.1301)

[acpi] Replace /proc/acpi/sleep

- Bad to remove proc file now, even though it's nearly useless. Reinstated in the name of compatibility.
- Restored original semantics - if software_suspend() is enabled, then just call that (and never go into low-power state). Otherwise, call acpi_suspend().
- acpi_suspend() is simply a wrapper for pm_suspend(), passing down the right argument. This is so we don't have to do everything manually anymore.
- Fixed long-standing bug by checking for "4b" in string written in to determine if we want to enter S4bios.

<mochel@osdl.org> (03/08/30 1.1300)

[swsusp] Restore software_suspend() call.

- Allows 'backdoor' interface to swsusp, as requested by Pavel.
- Simply a wrapper to pm_suspend(), though guaranteeing that swsusp is used, and system is shutdown (and put into low-power state).
- Call in sys_reboot() changed back to call to software_suspend().

<mochel@osdl.org> (03/08/30 1.1299)

[swsusp] Use BIO interface when reading from swap.

- bios are the preferred method for doing this type of stuff in 2.6. The __bread() uses bio's in the end anyway.
- bios make it really easy to implement write functionality, so we are able to reset the swap signature immediately after checking it during resume. So, if something happens while resuming, we will still have valid swap to use.
- Thanks to Jens for some help in getting it working several months ago.

<mochel@osdl.org> (03/08/29 1.1298)

[swsusp] Minor cleanups in read_suspend_image()

- Make resume_bdev global to file, so we don't have to pass it around (we always use the same one, so it shouldn't make a difference).

Linux–Kernel: Power Management Update

- Allocate cur in read_suspend_image(), since it's the only function that uses it.
- Check all errors and make sure we free cur if any happen.
- Make sure to return errors from the functions called, not our own.
- Free the pagedir if we hit an error after we allocate it.

<mochel@osdl.org> (03/08/27 1.1297)

[acpi] Move register save closer to call to enter sleep state.

- By moving acpi_{save,restore}_state_mem() into acpi_pm_enter(), implying after interrupts have been disabled and nothing else is running on the system, S3 is able to resume properly.

<mochel@osdl.org> (03/08/27 1.1296)

[power] Make sure devices get added to the PM lists before bus_add_device().

- Prevents ordering issues when drivers add more devices →probe().

<mochel@osdl.org> (03/08/26 1.1295)

[power] Separate suspend–to–disk from other suspend sequences.

- Put in kernel/power/disk.c
- Make compilation depend on CONFIG_SOFTWARE_SUSPEND (should probably be renamed to CONFIG_PM_STD or some such).

<mochel@osdl.org> (03/08/25 1.1294)

[power] Fix handling of pm_users.

- Actually decrement on device_pm_release()
- Call from device_pm_remove().

<mochel@osdl.org> (03/08/25 1.1292)

[power] Fix device suspend handling

- Handle –EAGAIN in device_suspend() properly: keep going, with error reset to 0.
- Call dpm_resume() if we got a real error, instead of device_resume(), which would deadlock.

<mochel@osdl.org> (03/08/22 1.1276.19.8)

[power] swsusp Cleanups

- do_magic()
- Rename to swsusp_arch_suspend().
- Move declaration to swsusp.c

Linux–Kernel: Power Management Update

- arch_prepare_suspend()
 - Return an int
 - Fix x86 version to return –EFAULT if cpu does not have pse, instead of calling panic().
 - Call from swsusp_save().

- do_magic_suspend_1()
 - Move body to pm_suspend_disk()
 - Remove.

- do_magic_suspend_2()
 - Rename to swsusp_suspend()
 - Move IRQ fiddling to suspend_save_image(), since that's the only call that needs it.
 - Return an int.

- do_magic_resume_1()
 - Move body to pm_resume().
 - Remove

- do_magic_resume_2()
 - Rename to swsusp_resume().
 - Return an int.

- swsusp general
 - Remove unnecessary includes.
 - Remove suspend_pagedir_lock, since it was only used to disable IRQs.
 - Change swsusp_{suspend,resume} return an int, so pm_suspend_disk() knows if anything failed.

<mochel@osdl.org> (03/08/22 1.1276.19.7)

[power] Move i386–specific swsusp code to arch/i386/power/

<mochel@osdl.org> (03/08/22 1.1276.19.6)

[power] Fix up sysfs state handling.

<mochel@osdl.org> (03/08/22 1.1276.19.5)

[power] Make sure console level is high when suspending.

<mochel@osdl.org> (03/08/22 1.1276.20.1)

[power] Fix sysfs state reporting.

–

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/>